How to Create Massively Scalable Database Applications

Doug Hood
@ScalableDBDoug
Consulting Member of Technical Staff
Product Manager TimesTen In-Memory Database
May 16, 2019

Agenda

1. Latency, Throughput and Scalability
2. Scale-up vs Scale-out
3. Scale-out Architectures
4. Trivial Scalability Benchmarks
5. Scaling a Customer Workload
6. Summary and Q & A

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

Latency, Throughput and Scalability

Latency	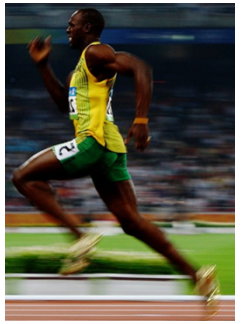	**How quickly can one operation complete** **One sprinter in 9.58 seconds** **~ 40 km/h for 100M [2009]**
Throughput		**How quickly can many operations complete** **Ten sprinters in under 11 seconds** **~ 40 km/h for 100m [2009]**
Scalability		**By adding more resources can throughput keep increasing** **33 cars on 2.5 mile oval track** **~250 km/h for 804 km [Indy 500, 2017]**

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

You can only go so big

IBM z14
z14 Microprocessor
32 CPUs @ 5.2 GHz
- 10 cores, 20 threads
8 TB DRAM

Oracle Super Cluster
SPARC M8
16 CPUs @ 5.1 GHz
- 32 cores, 256 threads
16 TB DRAM

HPE Superdome Flex
Intel Xeon
32 CPUs @ 3.6 GHz
- 28 cores, 56 threads
48 TB DRAM

SGI Altix 4700
Intel Itanium 2
2048 CPUs @ 900 MHz
- 2 cores, 4 threads
128 TB DRAM

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

Vertical Scaling Limits

- Only so many CPUs interconnected
- NUMA limits
- Complexity & Cost
- Niche Market

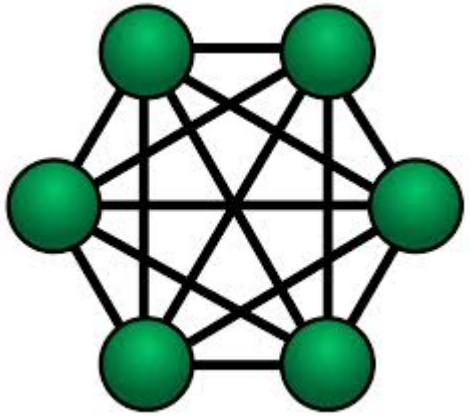

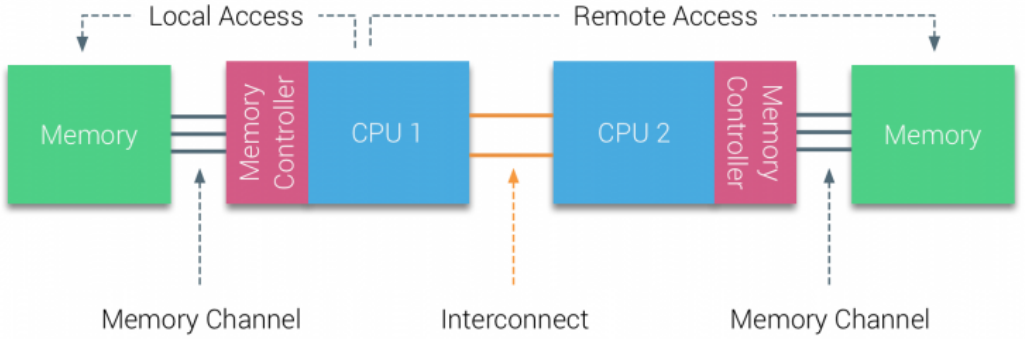

8+ Sockets

4-8 Sockets

1-2 Sockets

ORACLE

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

Horizontal Scaling hardware

- Use **cheap/fast Linux x8664 servers**, eg Oracle Sun X7-2
- NUMA affects are minimal
- Commodity servers keep getting **faster**, **cheaper** and **more powerful**
- 1.5 TB DRAM [Persistent Ram coming, **Intel/Oracle PMem demo**]
- Two Intel Xeon 8164 @ 2.2 GHz, 26 cores
- Up to eight NVMe SSDs
- **42 1U servers per Rack:**
  - 2*42 = **84 CPUs**
  - 1.5 * 42 = **63 TB RAM**

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

Lower Latency with TimesTen Cache

Query	Oracle	Cache
Q1	43	3
Q2	69	6
Q3	105	8
Q4	121	20
Q5	140	18
Q6	163	19
Q7	231	18

Oracle 11.2.0.4 RAC
RAC nodes were Oracle Sun X7-2L
NVMe Storage
Over 50 Million Users

Application Tier Database Cache (TimesTen)
Ran on the same nodes as the production RAC
5 table joins for 100s of millions of rows of data

Latency is in Micro Seconds …

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Oracle Database & Real Application Clusters Architecture

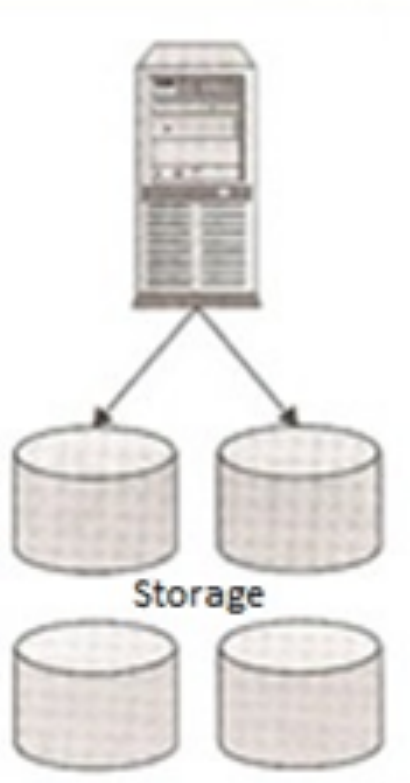

Oracle Database
- Single Instance
- Single DB image

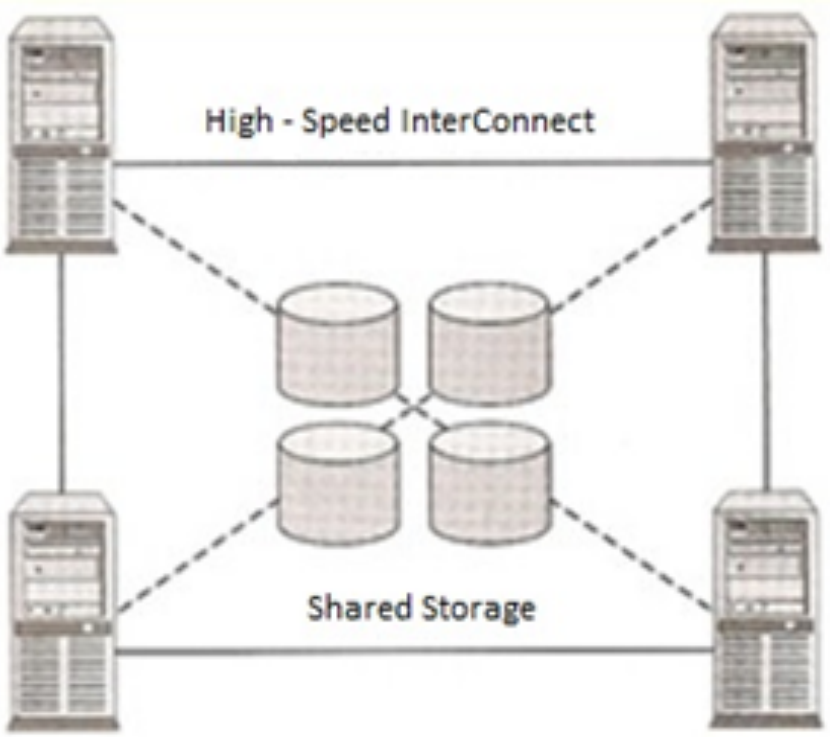

Oracle Real Application Clusters
- Multiple Database Instances
- Single DB image
- Shared Storage

Oracle Exadata
- Multiple Database Instances
- Single DB image
- Shared Storage

ORACLE

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

Oracle Sharding Architecture

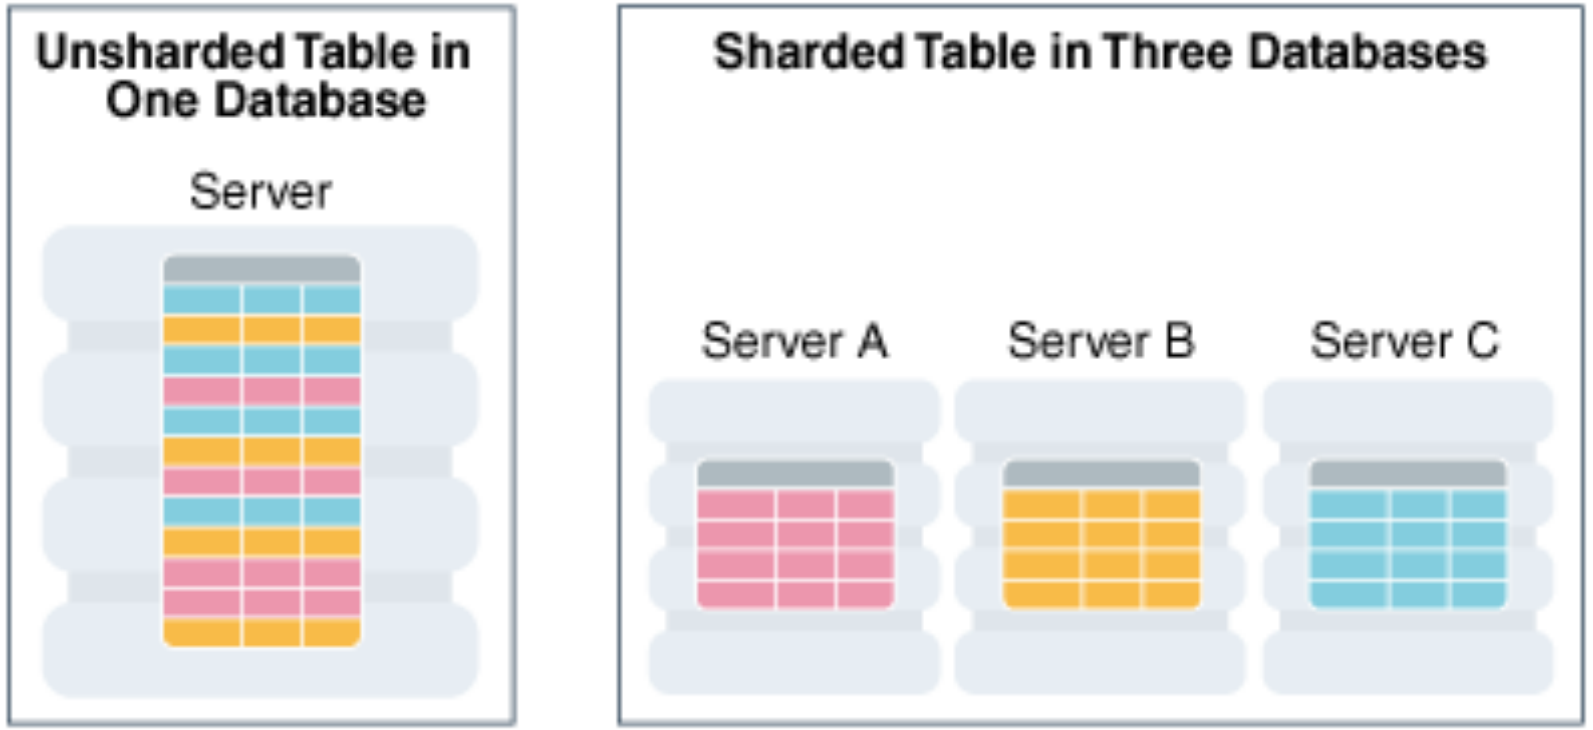

Oracle Sharding

- Multiple Database Instances
- Multiple DB images
- Independent Storage

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

Oracle NoSQL Architecture

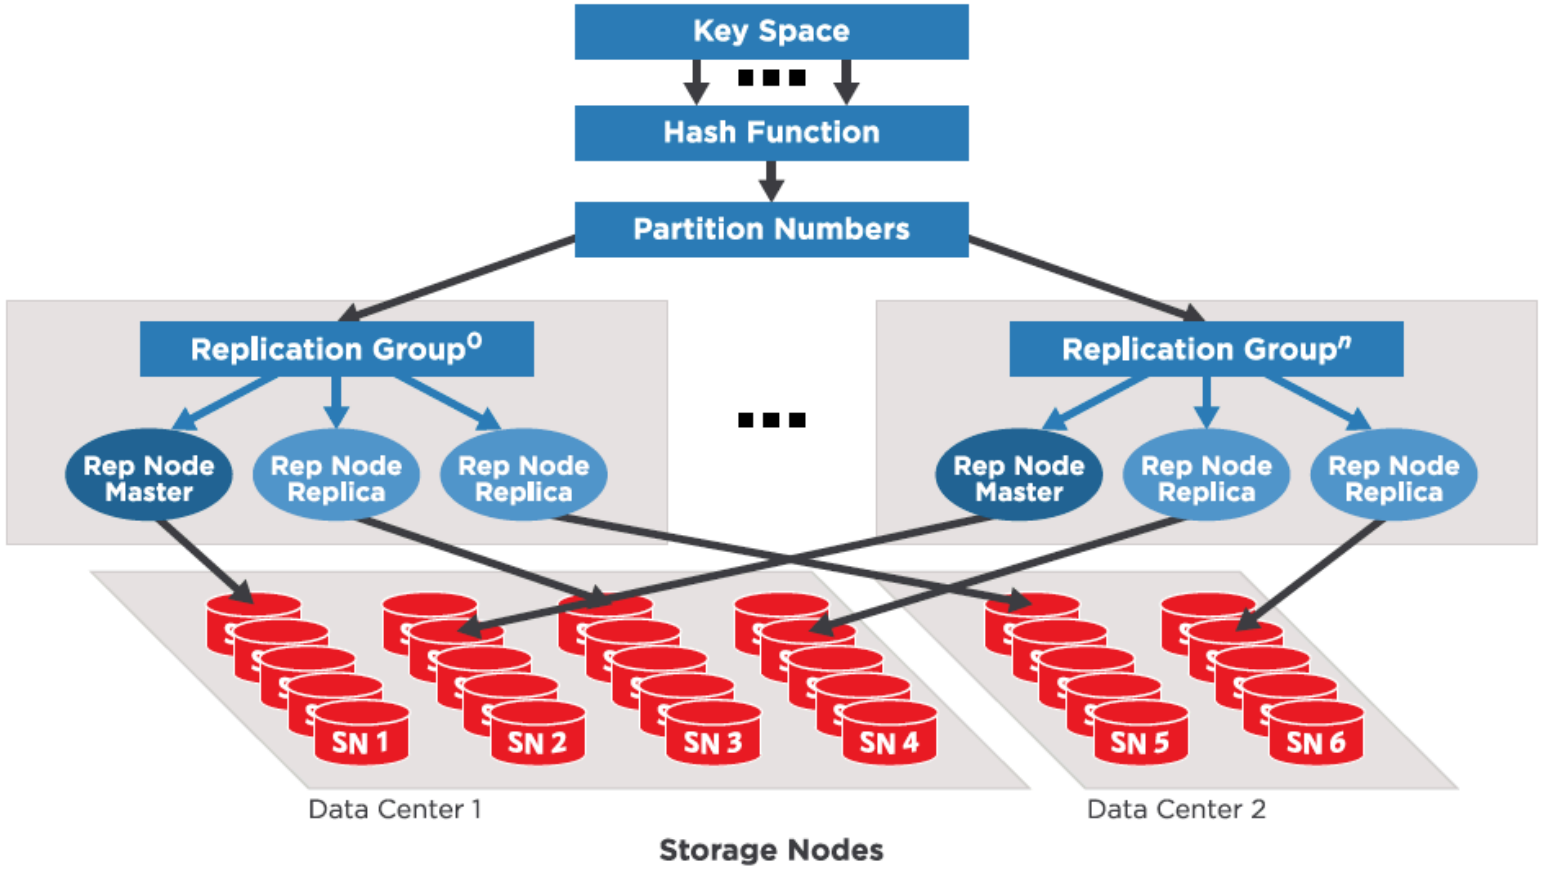

Oracle NoSQL

- Multiple ‘DB’ Instances
- One DB image
- Independent Storage

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

Oracle TimesTen Scaleout Architecture

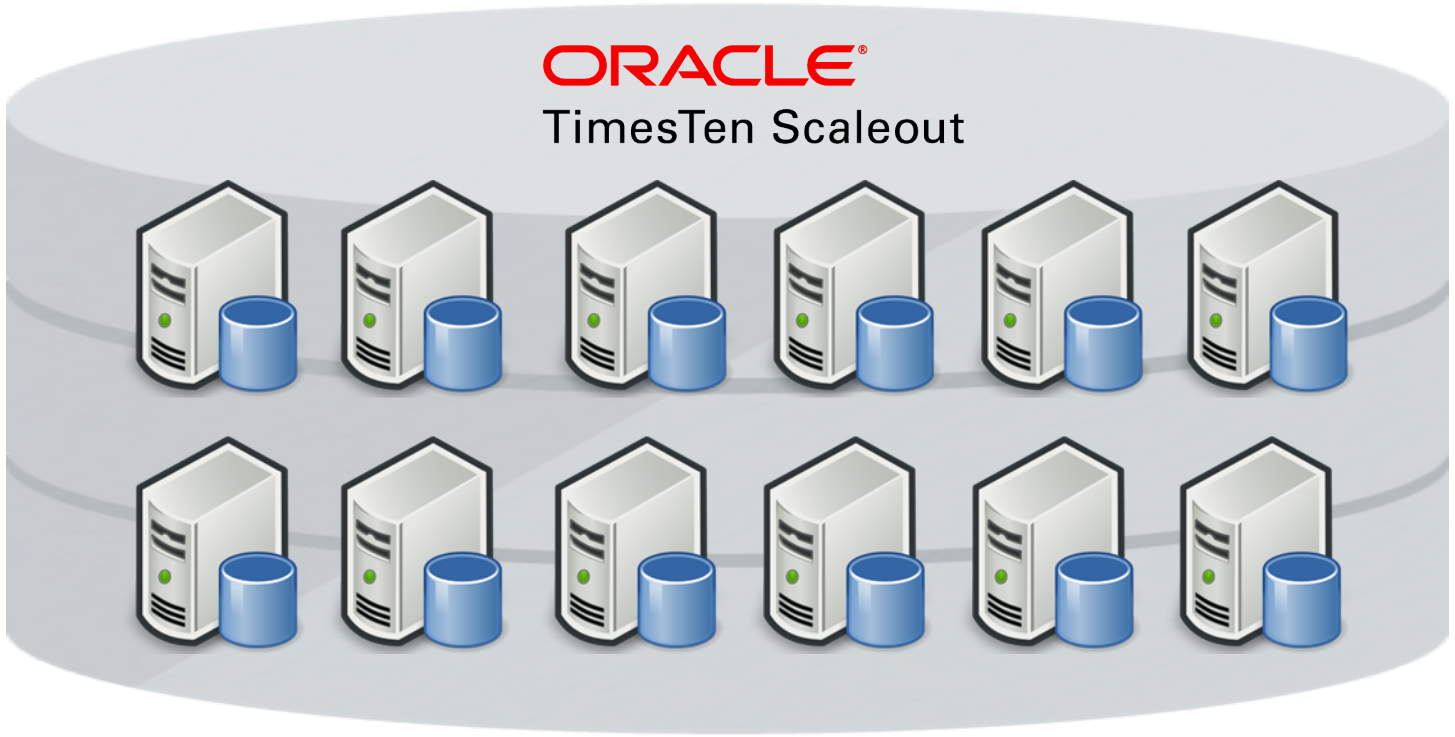

Oracle TimesTen Scaleout Architecture

- Multiple Database Instances
- Single DB image
- Shared Nothing

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

Summary of how to Scale Database Apps

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Summary of how to Scale Database Apps

- Do not do dumb things
- Tune your SQL
- Use PLSQL stored procedures intelligently
- Use good hardware
- Scale-up with Sun SuperCluster
- Scale-out with Exadata
- Scale-out with Application Tier Database Cache or TimesTen Scaleout

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Low Latency - Microseconds Response Time

```
select directory_nb,
   last_calling_party,
   descr
from vpn_users
where vpn_id = :1
and vpn_nb= :2
```

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)

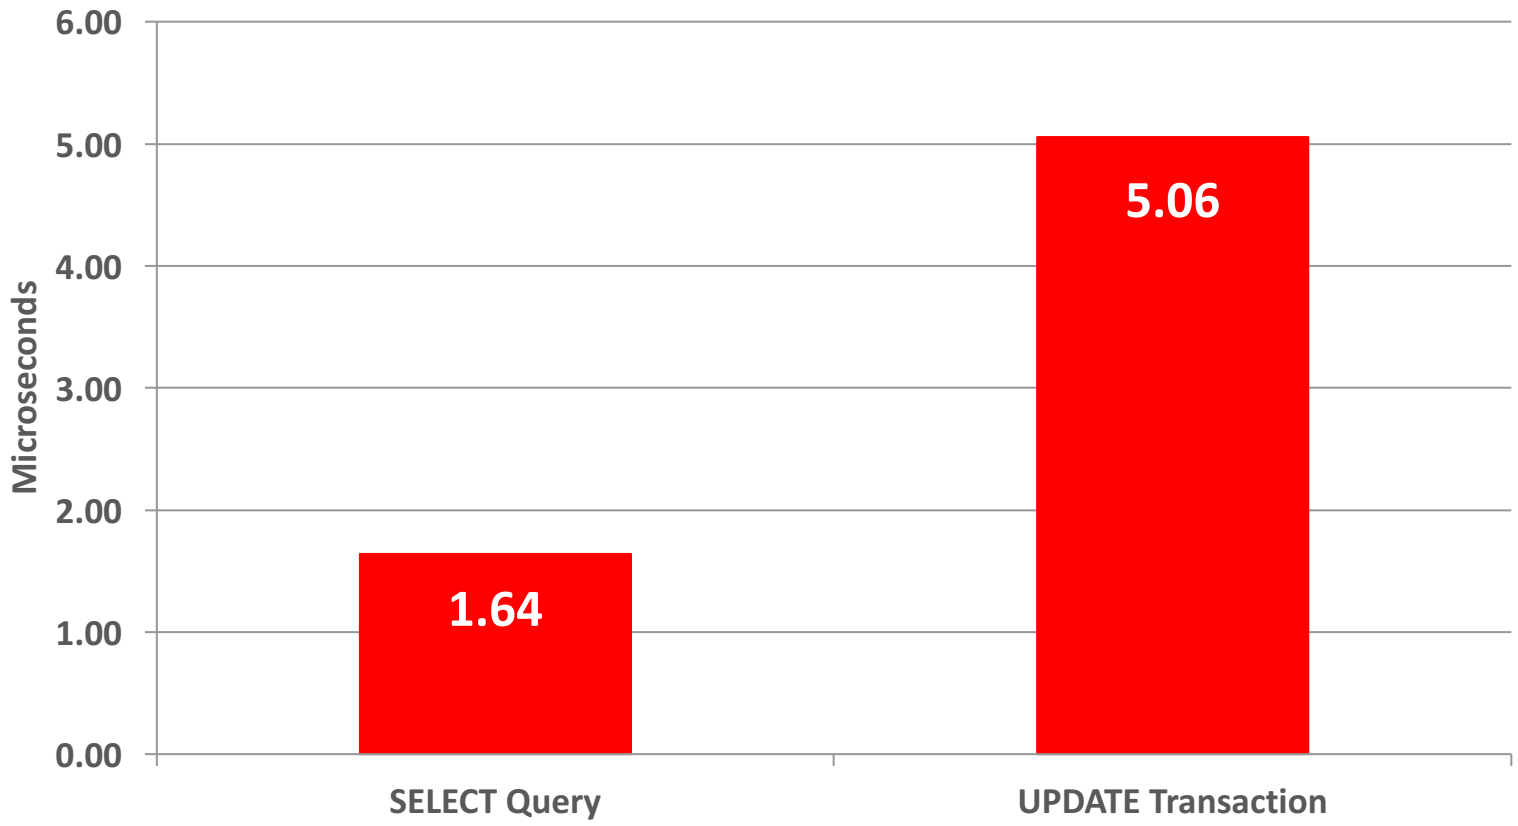

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. |

Some Throughput & Scalability Benchmarks

- 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

Surveyed YCSB (Workload B) Results*

Product	Type	Nodes	Ops/Sec
cassandra	NoSQL DB	32	**227 K**
mongoDB	NoSQL DB	2	**275 K**
SCYLLA.	NoSQL DB	3	**715 K**
VOLTDB	Scale-Out RDBMS	6	**1.6 M**
AEROSPIKE	NoSQL DB	8	**1.6 M**

** There is no official repository of YCSB results*
These were the largest results we found online

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

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

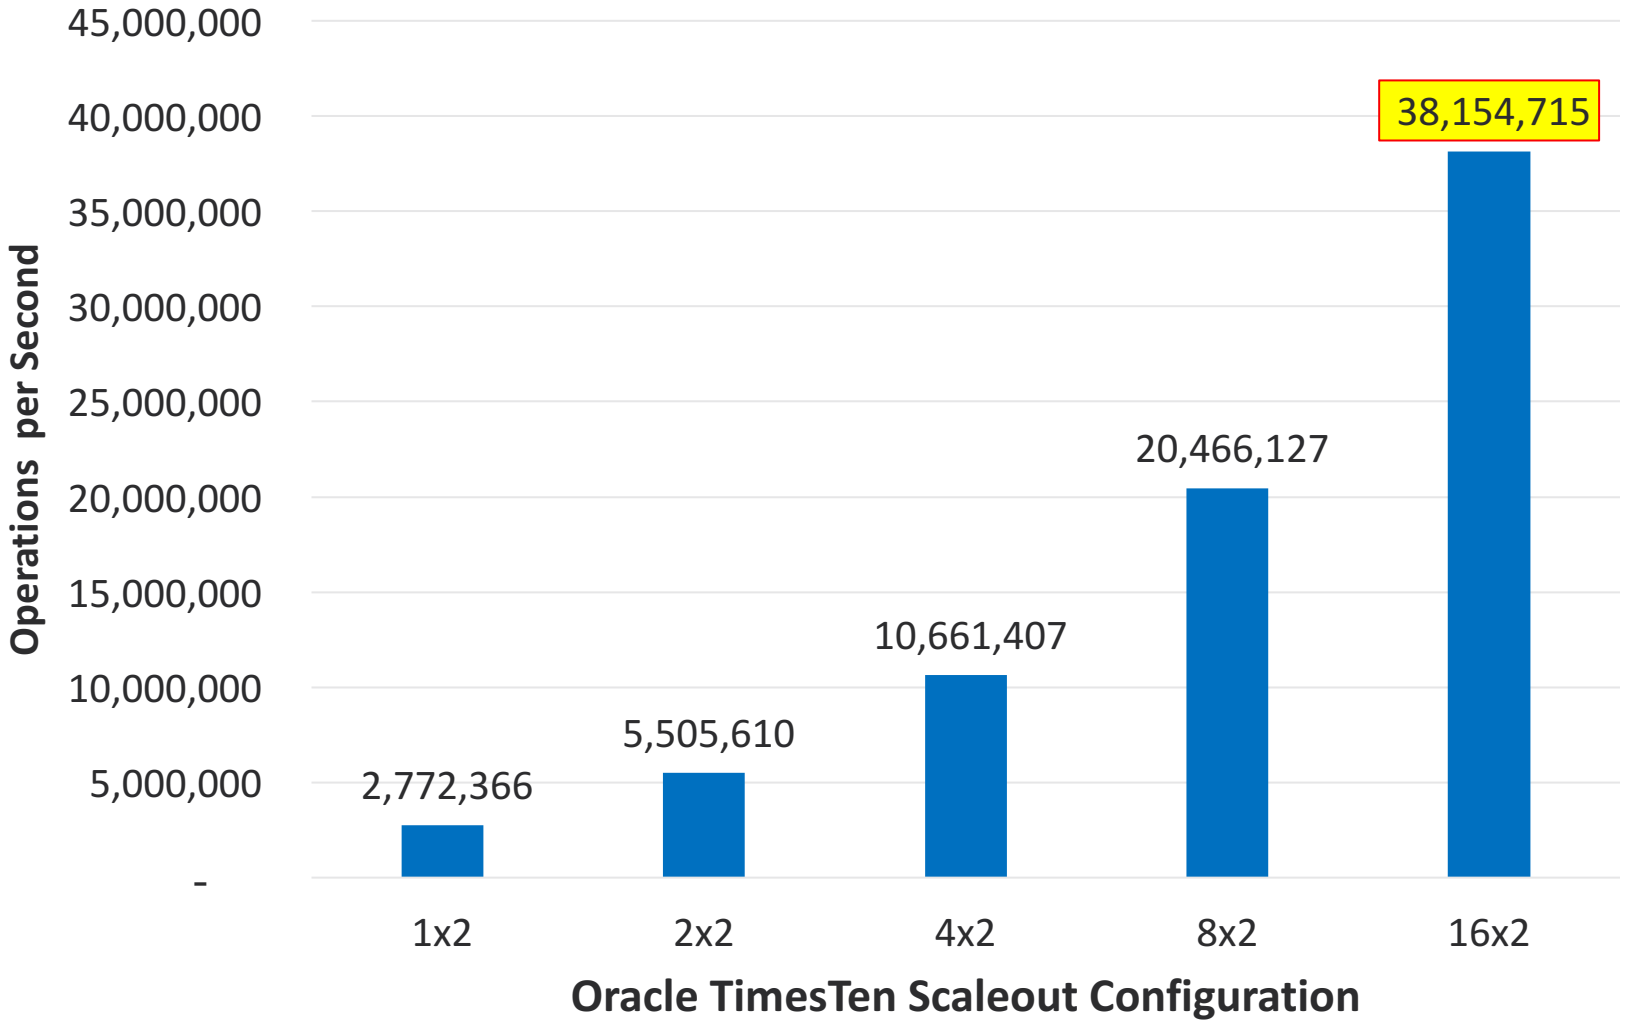

YCSB version 0.15.0

- 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

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

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

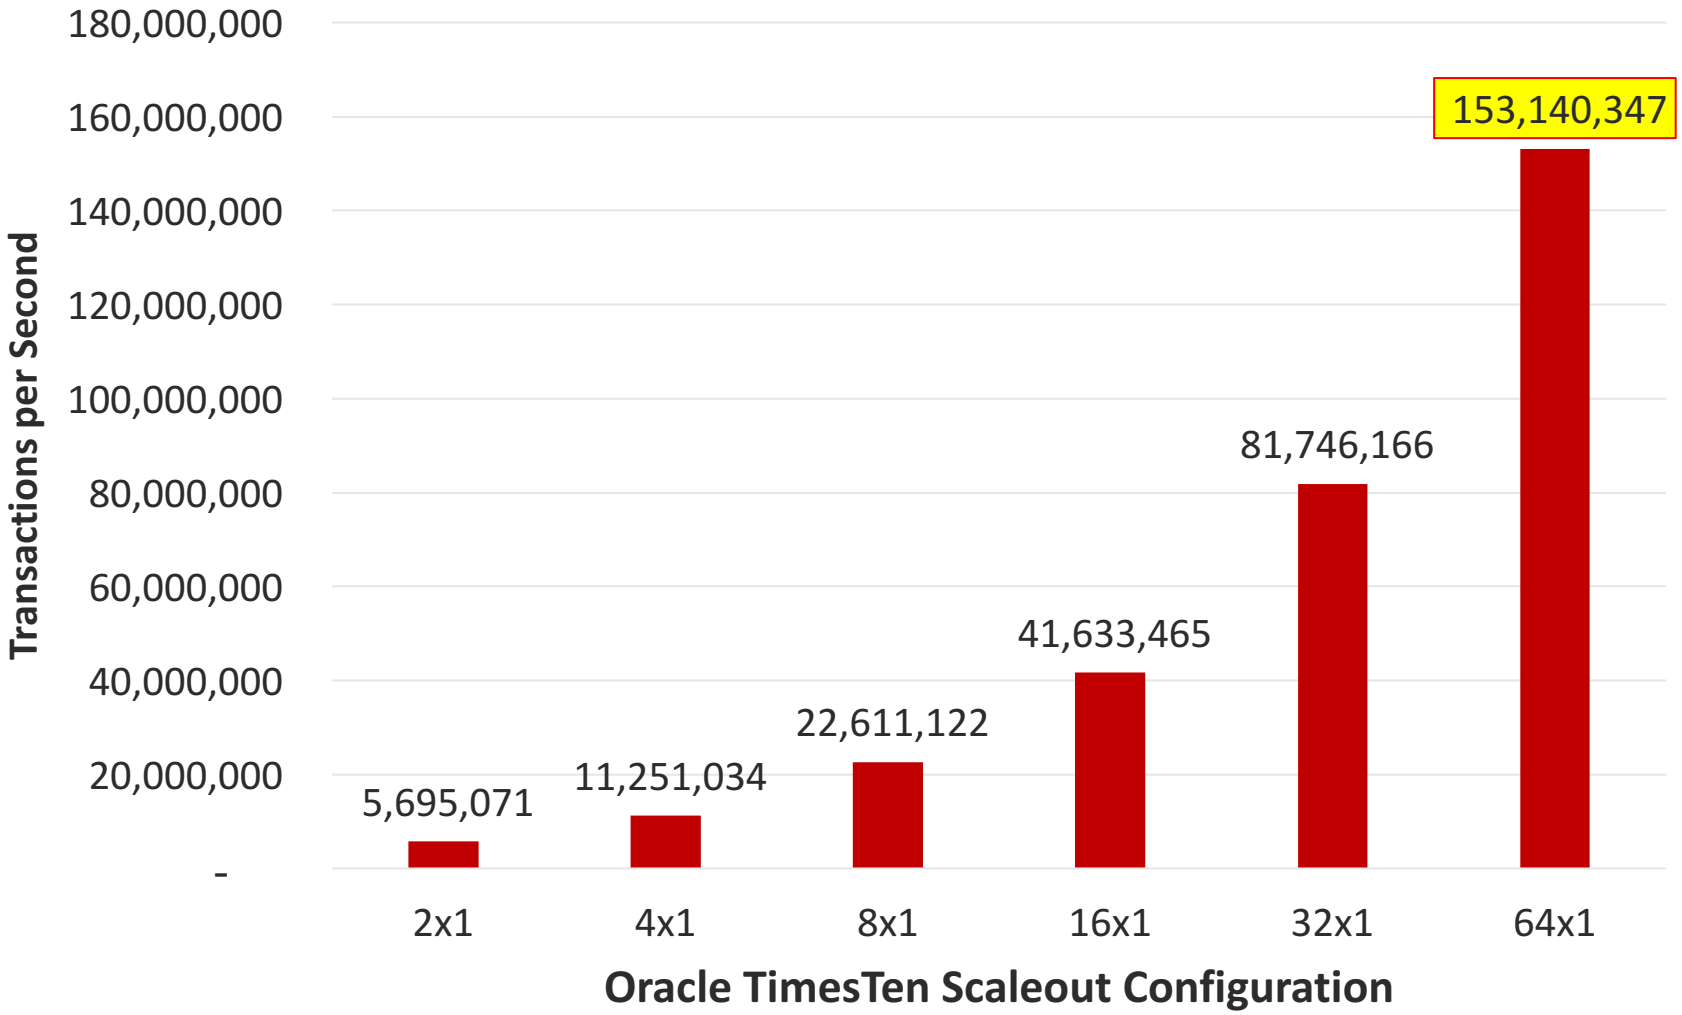

TPTBM Configuration

- 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

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

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

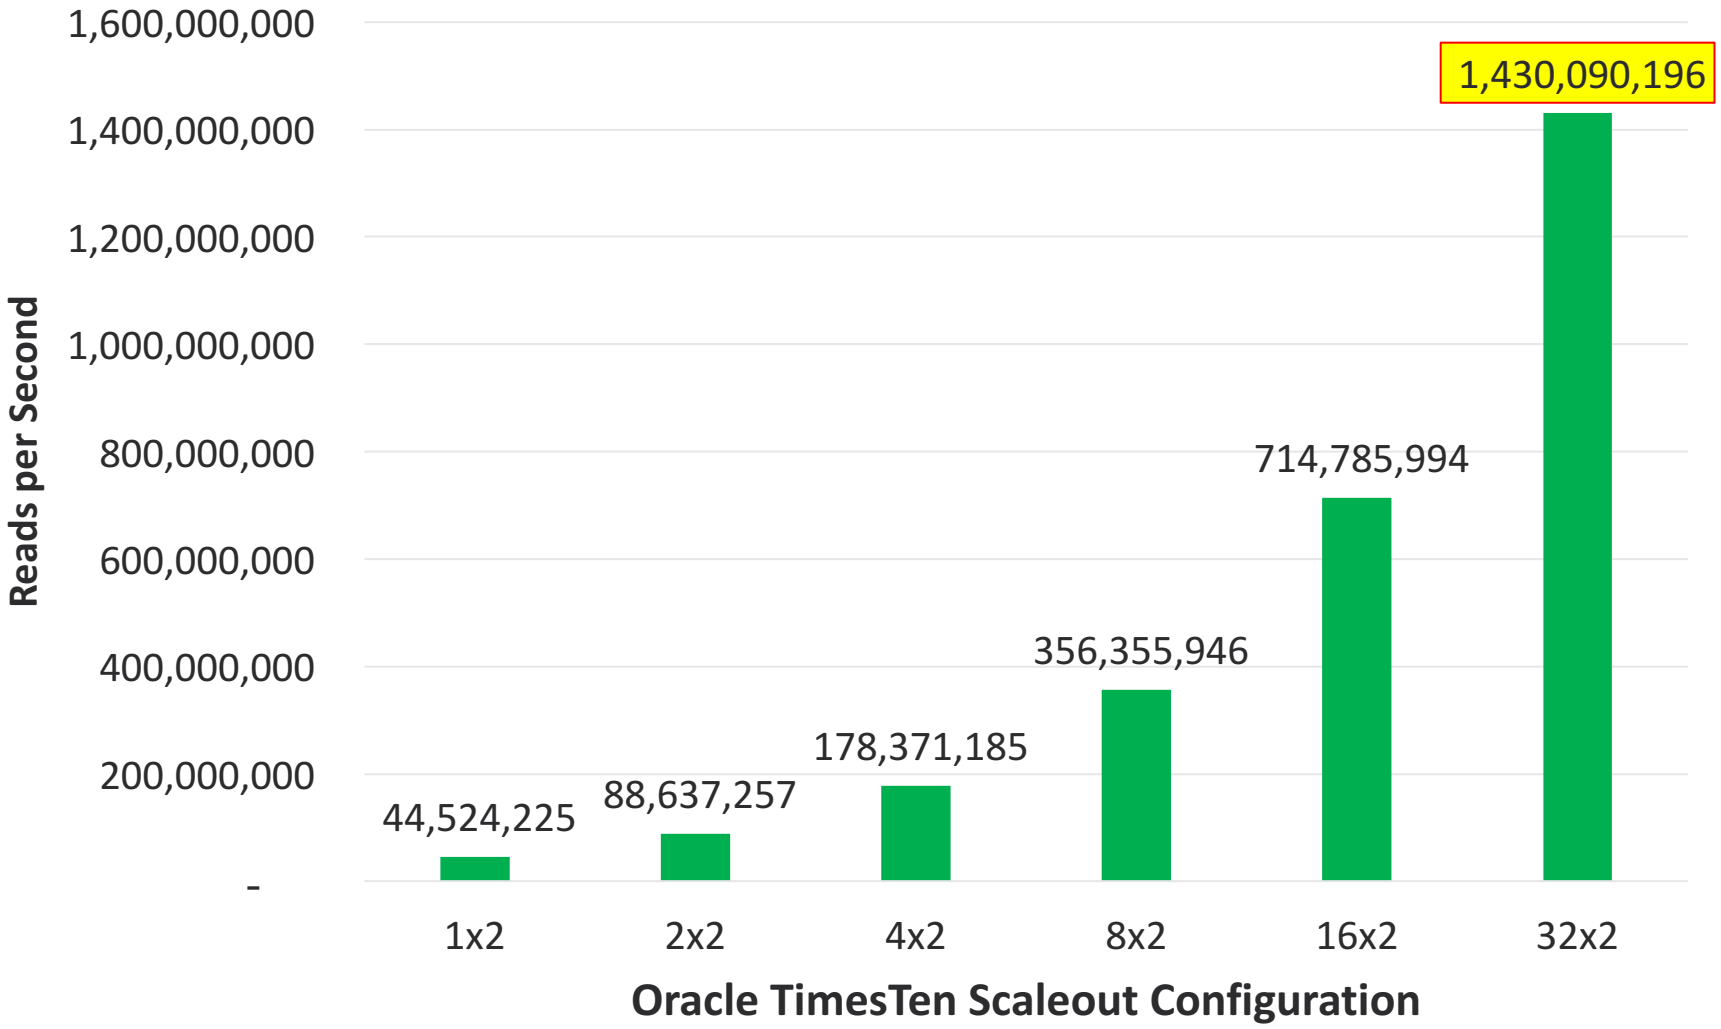

TPTBM Configuration

- 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

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

What Hardware was Used?

Oracle Sun X7-2

- Two Intel Xeon 8164 @ 2.2 GHz, 26 cores
- 768 GB RAM
- Four NVMe SSDs
- Two 10G Ethernet

Oracle Cloud Infrastructure

- 32 * BM.DenseIO2.52

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

ORACLE®

TimesTen Scaleout

World's Fastest OLTP Database

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Subset of Customer's Data Model

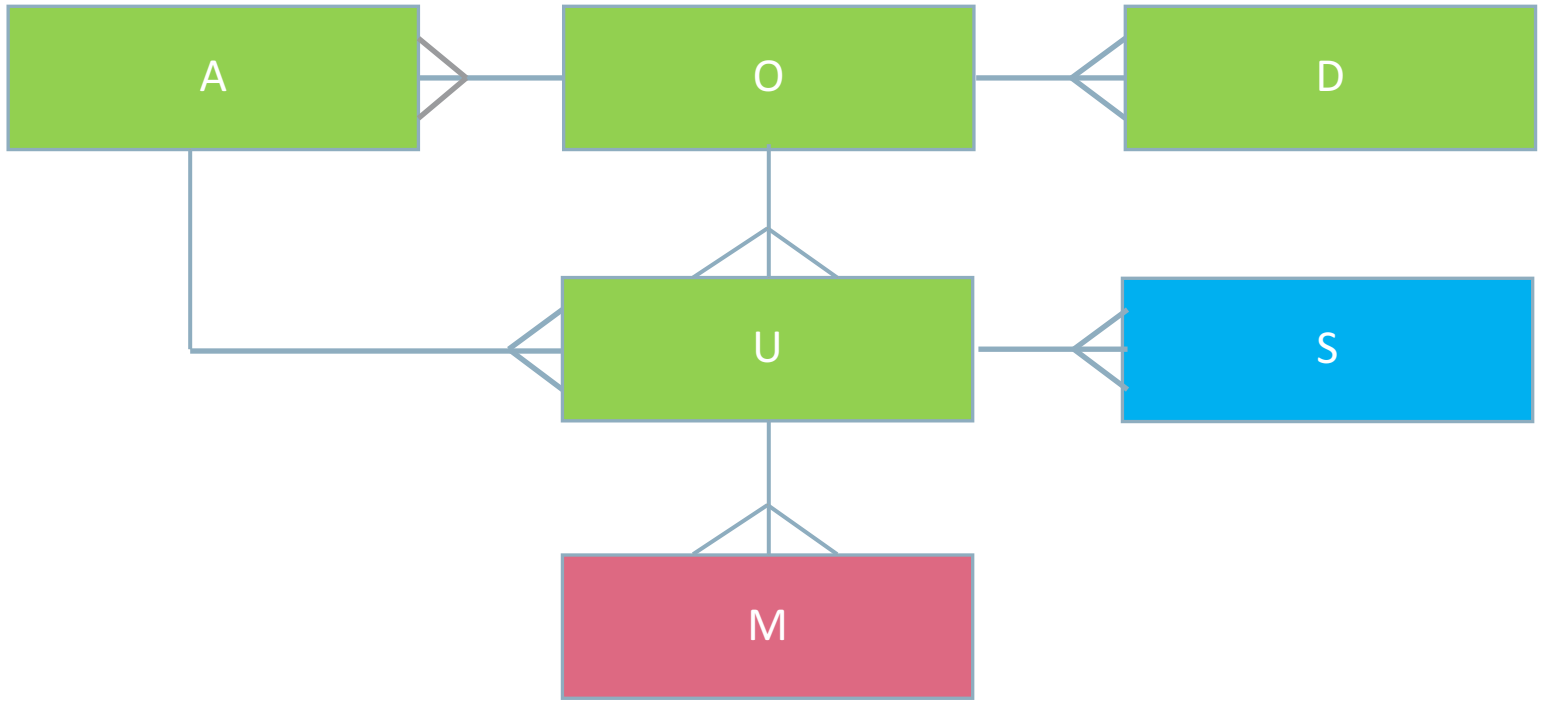

+ seven other tables for the 'write' workload

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

Critical Query

```
SELECT a.usr_id, …
FROM u, d, o, a
WHERE u.login_name = :loginName
AND u.dom_id = a.dom_id
AND u.usr_org_id = o.org_id
AND u.account_id = a.acct_id (+)
AND u.status <> :x;
```

```
SELECT mn_usr_id
FROM m
WHERE  mn_usr_id = uid
AND status = :y;
```

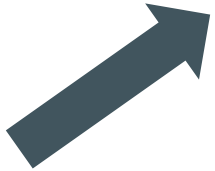

```
SELECT s.attr_name
FROM s
WHERE  s.entity_id = muid
AND (s.context = :p or b.context = :q)
AND (s.spid = :m or
        s.spid = :n or
        s.sid = :o)
ORDER BY b.attr_name;
```

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Oracle Confidential – Internal/Restricted/Highly Restricted

Critical Update Transaction

```
select something
from R1
where col1  = :x
and col2  = :y;

update  R1
set something  = :s
where col1  = :x
and col2  = :y;
```

```
select something
from R2
where col1  = :x
and col2  = :y;

update  R2
set something  = :s
where col1  = :x
and col2  = :y;
```

```
select something
from R3
where col1  = :x
and col2  = :y;

update  R3
set something  = :s
where col1  = :x
and col2  = :y;
```

- - - - - - - - - - - -

```
select something
from R7
where col1  = :x
and col2  = :y;

update  R7
set something  = :s
where col1  = :x
and col2  = :y;
```

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Scale Up or Scale Out?

Four 5.1 GHz SPARC CPUs
256 hardware threads per CPU socket
64 MB L3 Cache
16 TB RAM
8 NVMe SSD for DB Storage + 12 Disks
40 G Infiniband
4 Quad 10G Ethernet

Oracle Database 11g

cassandra
MEMSQL
mongoDB
AEROSPIKE
Couchbase

cassandra
MEMSQL
ORACLE
TimesTen Scaleout

KVM
CentOS
openstack
CLOUD SOFTWARE
NetApp

32 Core VMs
64 GB RAM
Cinder Storage

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Best Case Architecture for customer workload

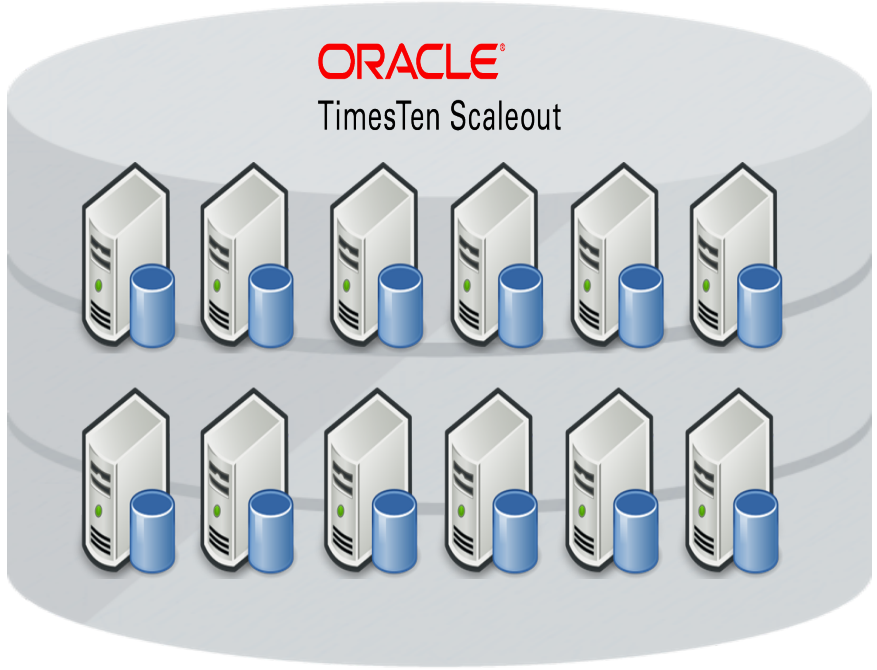

Attribute	Value
Data Reads & Writes	100% Local RAM
Storage Reads & Writes	100% Local NVMe SSD
Storage Bottleneck	No
Fast CPU	Xeon
Number of CPU cores	24
Sufficient Memory	Yes. 320 GB
DB Tuned	Yes
App tuned	**No. Python without SQL prepares or binds**

Result: 11 Million Transactions / second

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Worst Case Architecture for customer workload

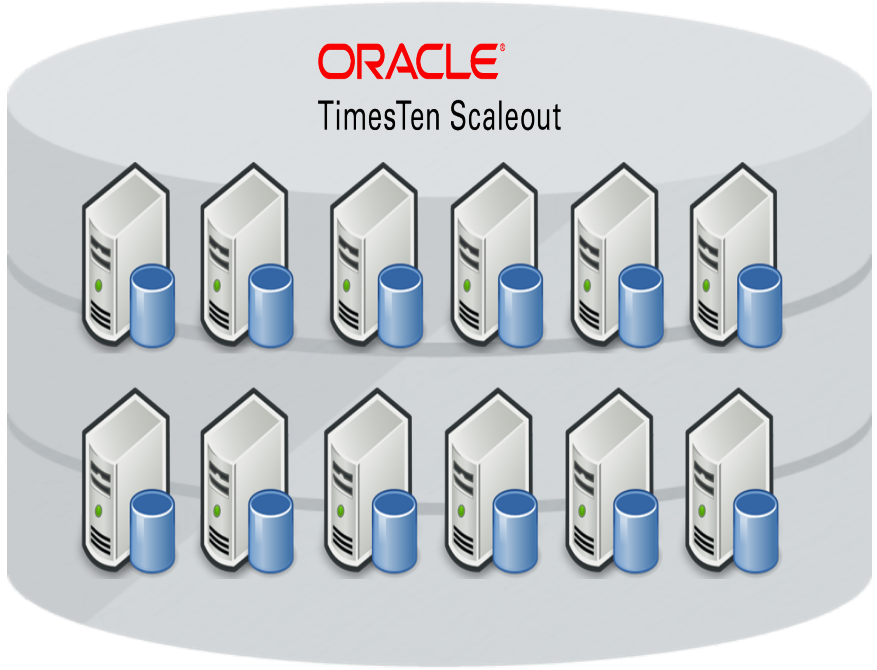

Attribute	Value
Data Reads & Writes	90% on a remote VM
Storage Reads & Writes	100% remote [Cinder/Netapp]
Storage Bottleneck	Maybe. Network bound
Fast CPU	Xeon
Number of CPU cores	32
Sufficient Memory	No. Only 32 GB
DB Tuned	Yes
App tuned	Yes. ODBC with SQL prepares and binds

Result: 304K Transactions / second

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Some Results

240K TPS
60/40 Workload
IO Bound
ACID 1PC

4 Socket SMP

cassandra

< 168K TPS
60/40 Workload
Network Bound
Eventual Cons

Negative Scaling

MEMSQL

168K TPS
60/40 Workload
Network Bound
Eventual Cons

37 Node Cluster

ORACLE®
TimesTen Scaleout

304K TPS
60/40 Workload
Network Bound
ACID 2PC

10 Node Cluster

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

How Many Client Server SQL Network Round Trips ?

1. **Select * from table where PK = :value;**
2. **Select * from table where PK between 10 and 20;**
3. **Update table set column = :X where PK = :value;**
4. **Update table set column = :X where PK between 1000 and 2000;**
5. **Select * from a, b, c, d where {non Cartesian Product}**

A. One
B. Two
C. Three
D. Lots
E. It Depends

How many server side network messages When tables are hash distributed?

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Data Distribution Methods

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
- Distribute by **Duplicate**
  - Identical copies on all elements
  - Useful for reference tables
  - Read and join optimization

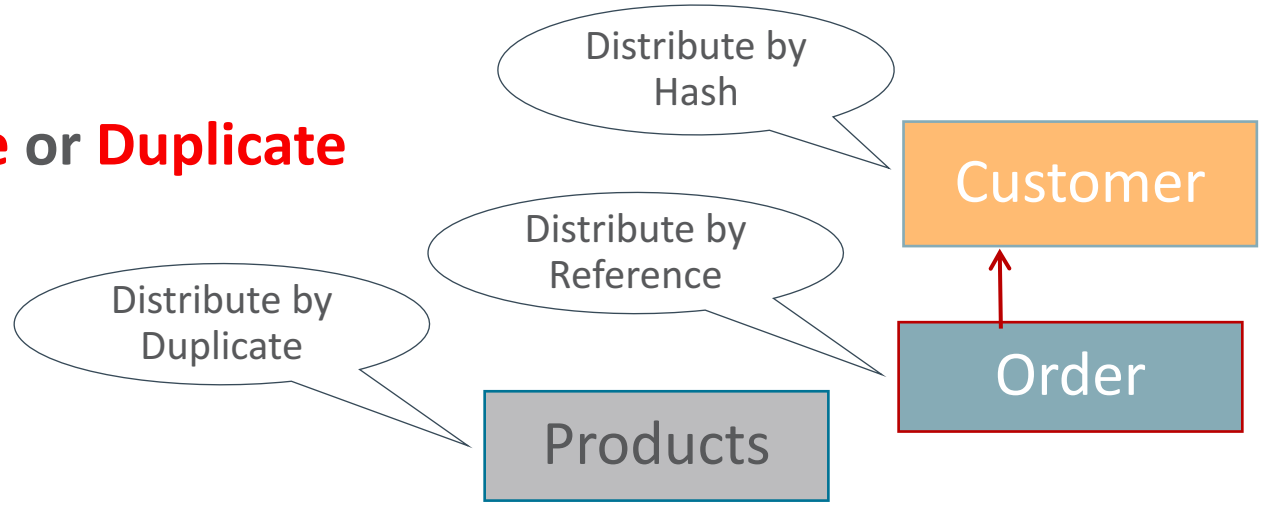

Element 1	
0	David
4	Igor
8	Tim

1	0	16/6/15
6	8	16/3/22

phone	100
tablet	200
watch	300

Element 2	
1	Bill
5	Sam
9	Charles

2	5	16/2/22

phone	100
tablet	200
watch	300

Element 3	
2	Olaf
6	Henri
10	Jie

5	6	16/5/10

phone	100
tablet	200
watch	300

Element 4	
3	Chi
7	Simon
11	Chris

3	3	16/3/1
4	11	16/2/5

phone	100
tablet	200
watch	300

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

Scalability Challenges

- Four table joins with hash distribution for ‘read workload’ with (+)
- Seven queries + seven updates for ‘write workload’
- Client Server round trips
- Not enough RAM [64 GB] per VM
- KVM + OpenStack Neutron networking overhead

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Techniques which helped Scalability

- Determine the best distribution clauses
  - The Distribution Advisor eliminates the guess work
- Determine the best indexes
  - The Index Advisor eliminates the guess work
- Prepare and Bind the SQL statements
- Check the explain plans
- Use Stored Procedures for the ‘read’ and ‘write’ transactions
  - Execute many statement in a single network round trip. Procedural logic + commit/rollback
- Use the Routing API
  - Determine where the data is to avoid network hops
- Use more DB nodes
  - The VNIC became network bound [ksoftirq]
  - Use more modes to lessen the load per VNIC

TODO
- TCP tuning
- RDMA

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

TimesTen Scaleout SQL APIs

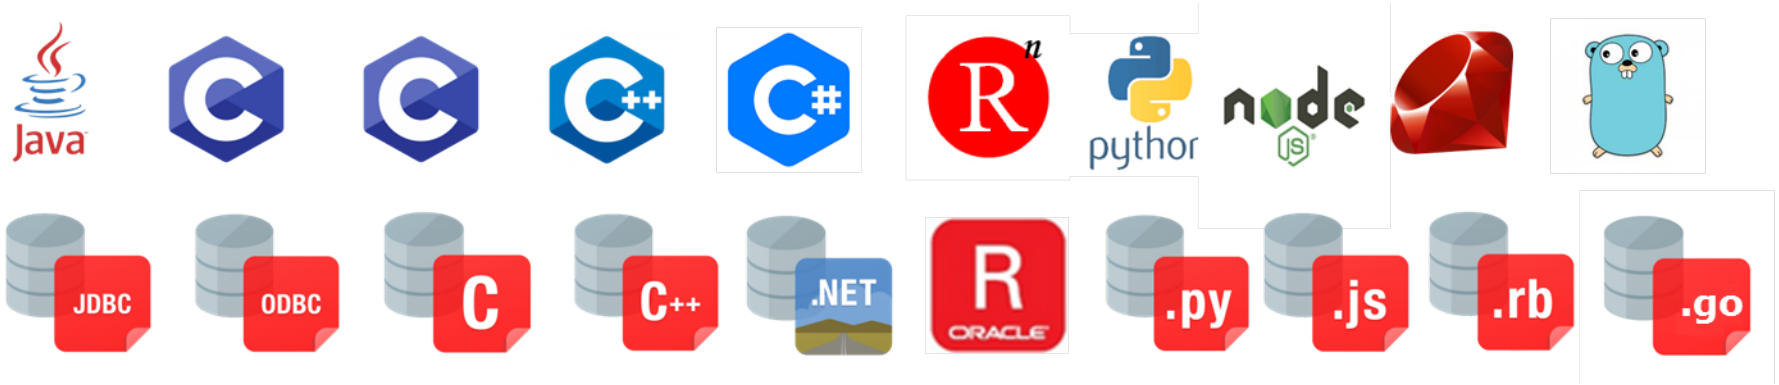

API	Comment
JDBC	The same (JDBC 4.3)
ODBC	The same (ODBC 3.5.2)
OCI	The same (OCI 11.2.0.4.+)
R-Oracle	The same (OCI 11.2.0.4.+)
ODP.Net	The same (OCI 11.2.0.4.+)
PL/SQL	The same (11.2.0.4.+)
Python	The same (cx_Oracle, ODPI-C)
Ruby	The same (Ruby-ODPI, ODPI-C)
GoLang	The same (go-goracle, ODPI-C)

ORACLE

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

TimesTen in On Premises

- TimesTen Scaleout requires :
  - Linux x8664 (glibc 2.12+)
    - Oracle Linux / Red Hat / CentOS 6.4+, 7+
    - Ubuntu 14.04+
    - SuSE 12+
  - JDK 8+
  - TCP/IP or IPoIB
  - A file system [eg ext4, not ext2 or ext3]
  - Enough RAM for the DB

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

TimesTen Scaleout on OCI, AWS, Azure, Google

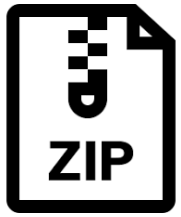

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. |

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

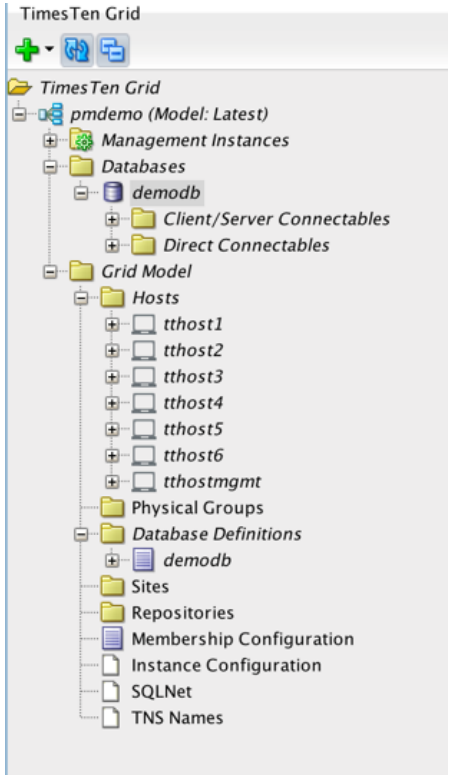

Status | Database Definition | Topology

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

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

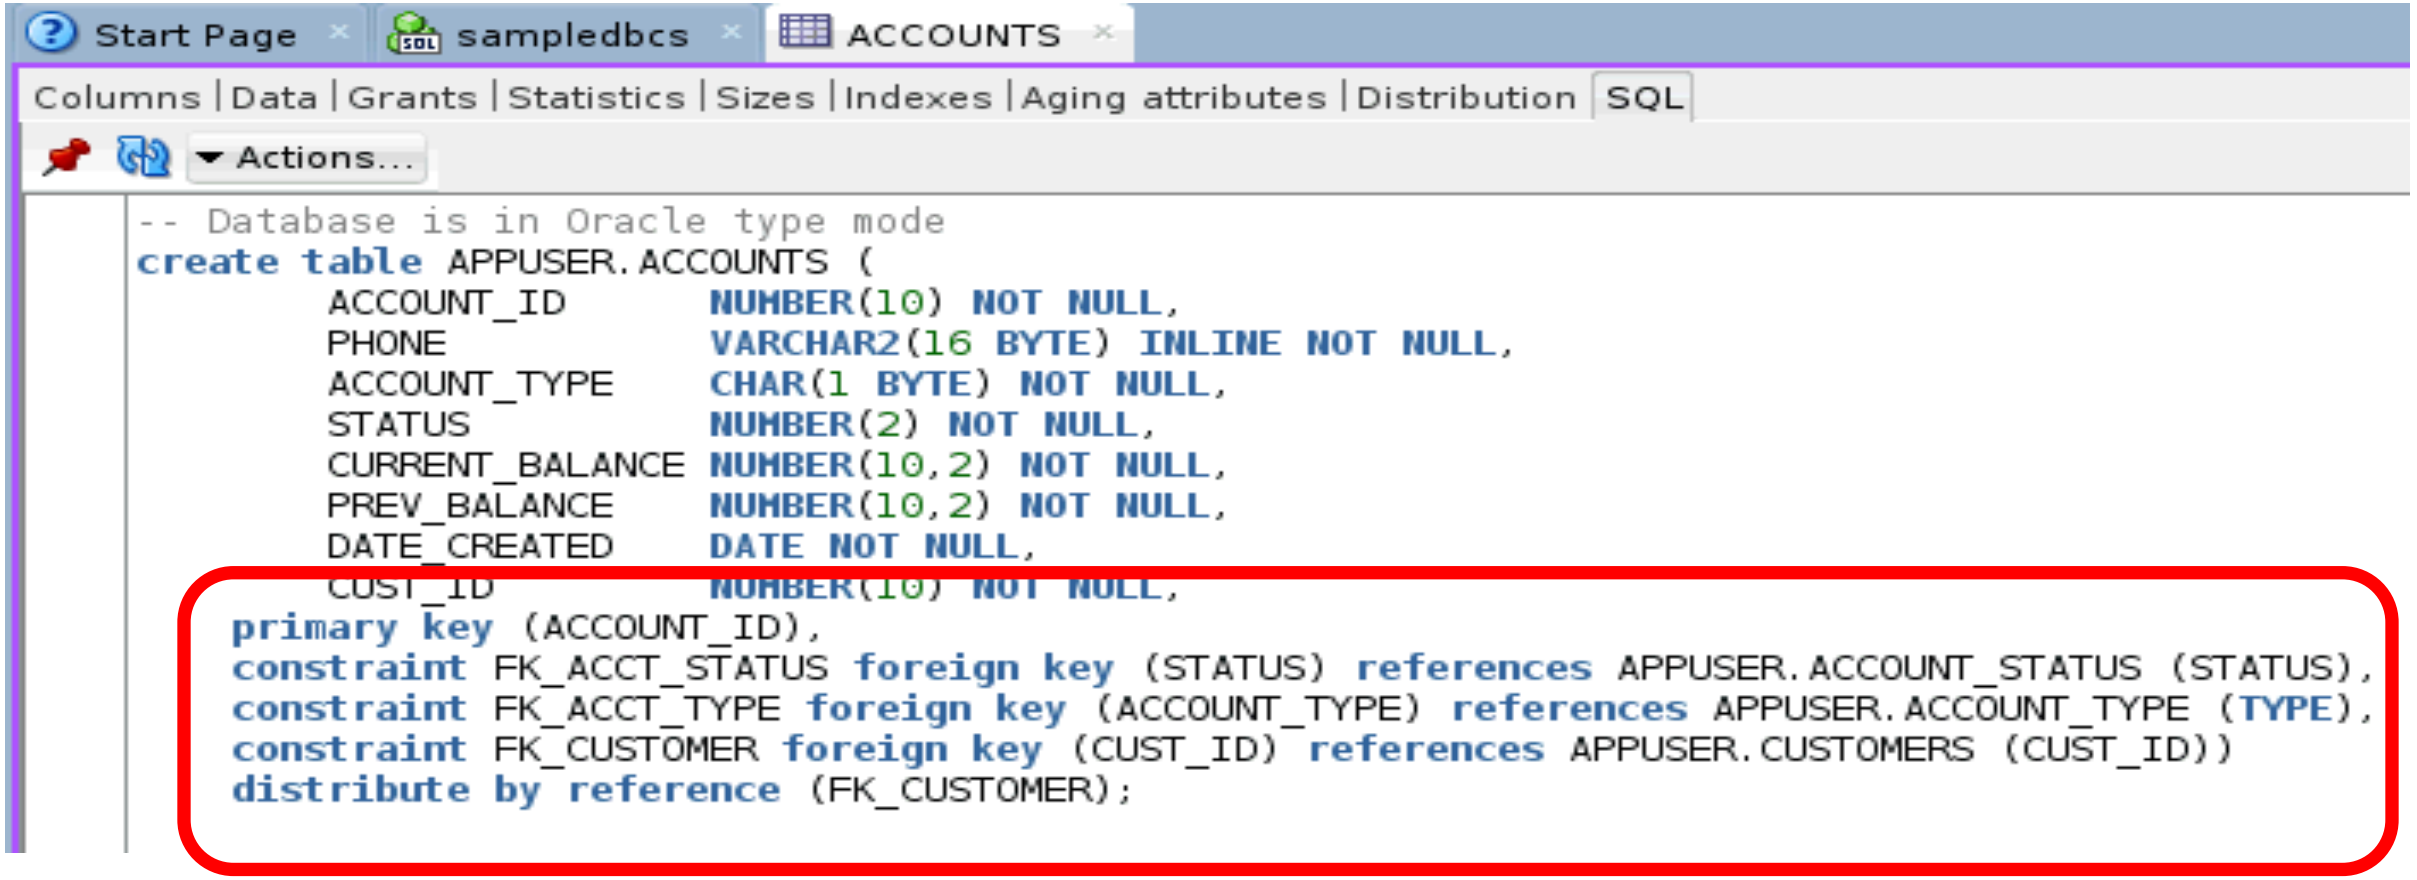

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

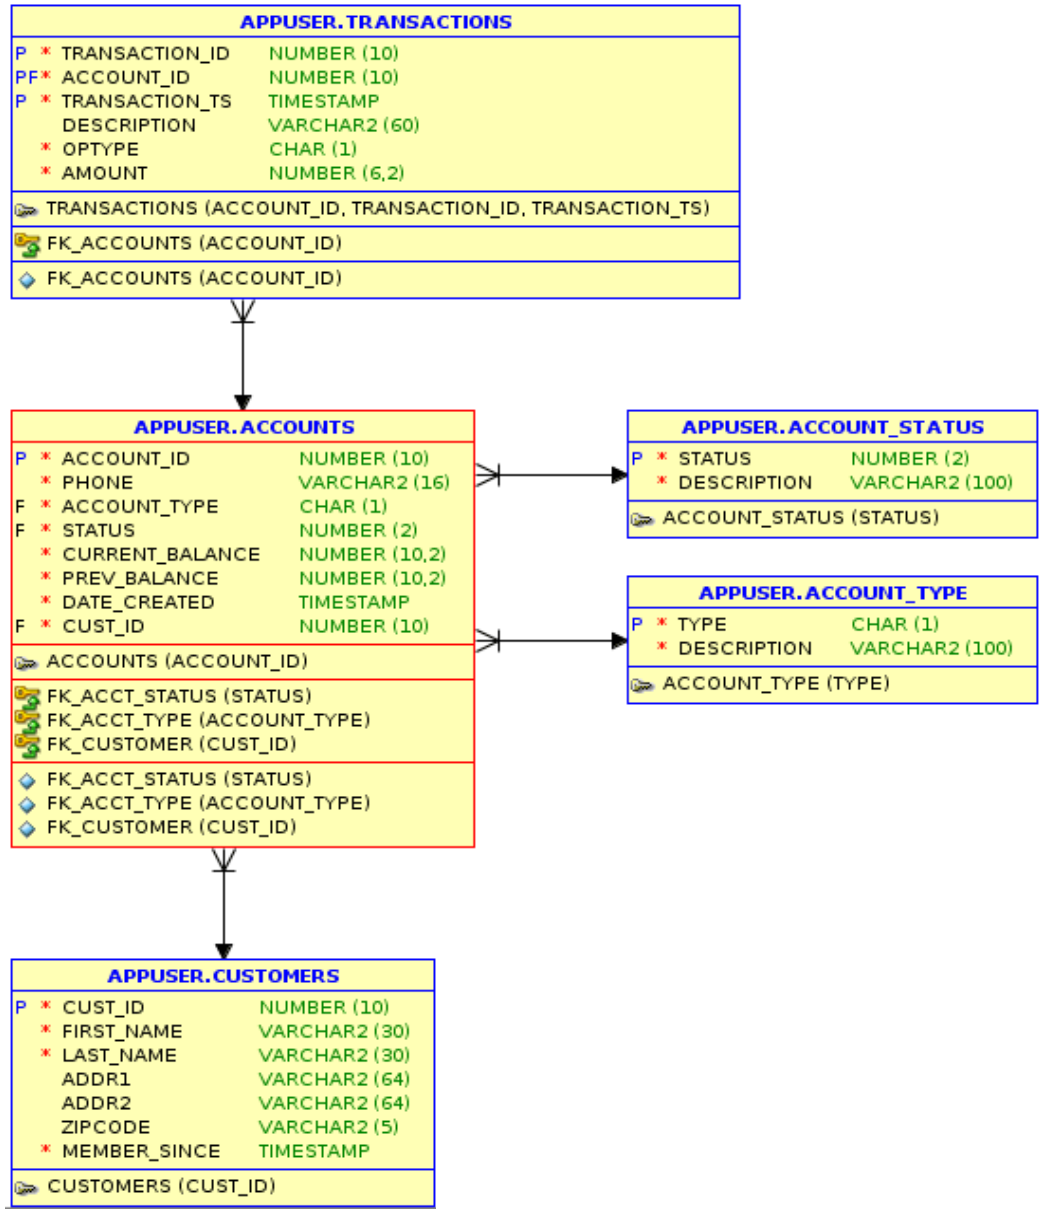

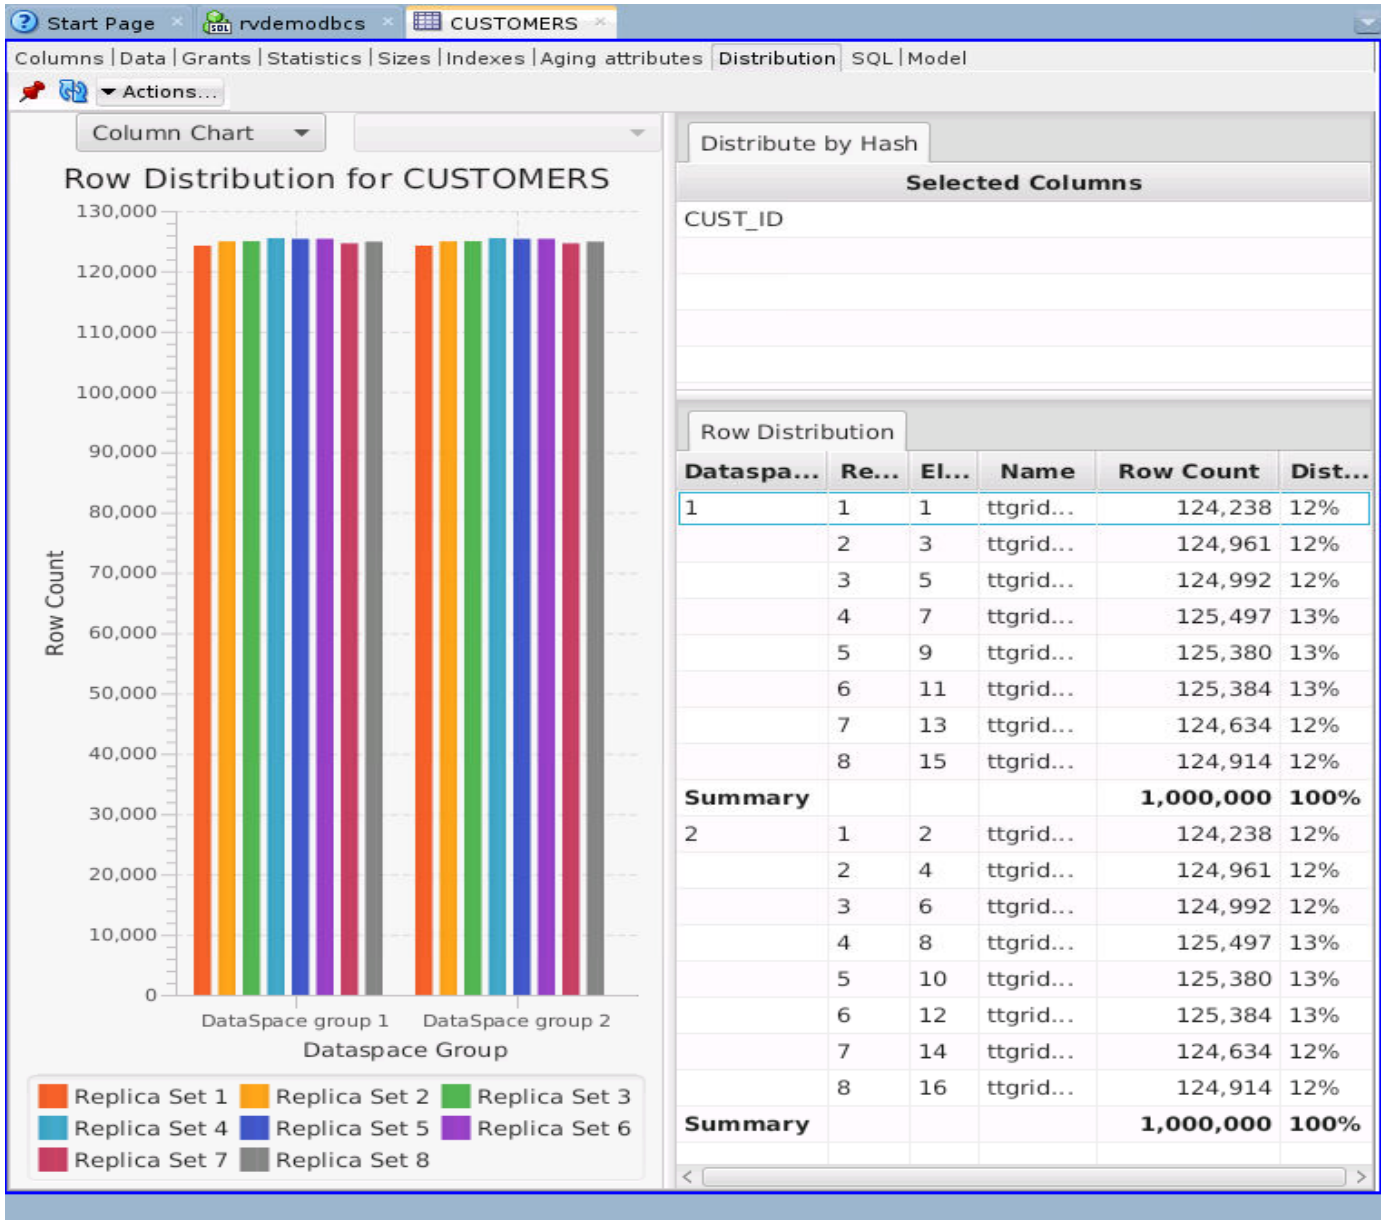

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Using Oracle cx_Python with TimesTen Scaleout

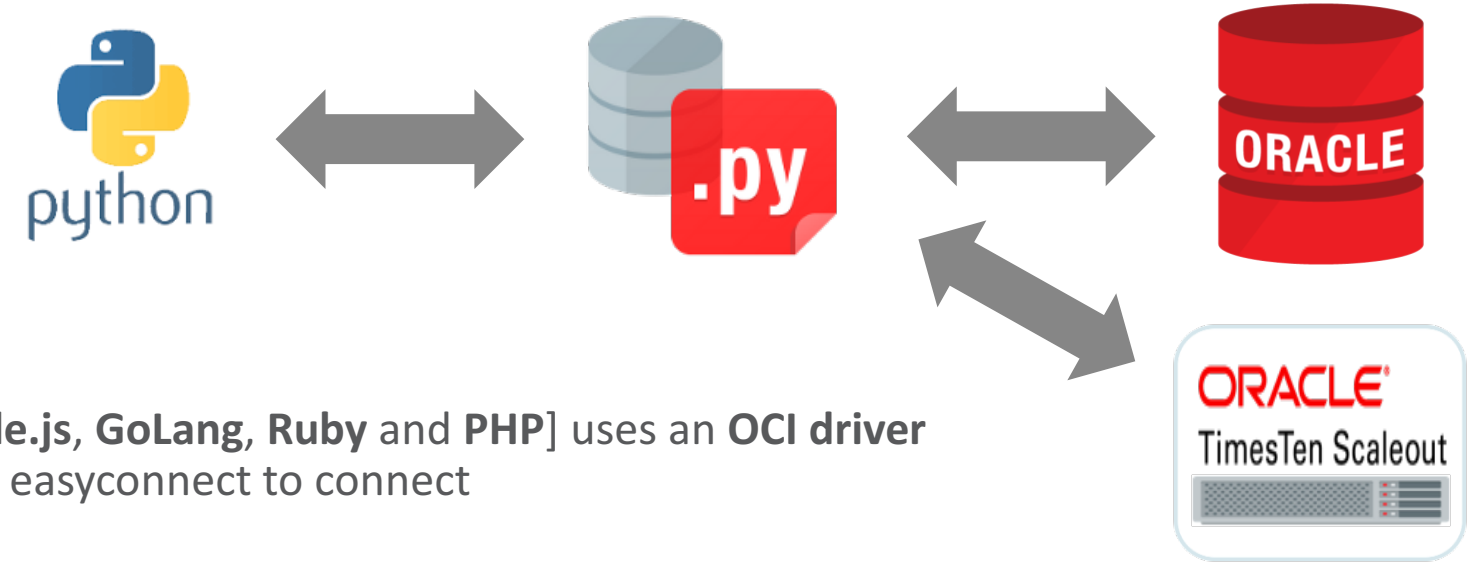

Python [and **Node.js**, **GoLang**, **Ruby** and **PHP**] uses an **OCI driver**
Use tnsnames or easyconnect to connect

tnsnames.ora :

sampledb_1812 =(DESCRIPTION=(CONNECT_DATA = (SERVICE_NAME = **sampledb_1812**)(SERVER = **timesten_direct**)))
sampledbCS_1812 =(DESCRIPTION=(CONNECT_DATA = (SERVICE_NAME = **sampledbCS_1812**)(SERVER = **timesten_client**)))

TimesTen ODBC DSN

Client/Server or Direct Linked

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. | Confidential – Oracle Internal/Restricted/Highly Restricted

Oracle TimesTen In-Memory Database

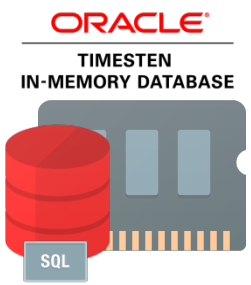

Relational Database

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

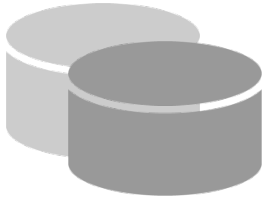

Persistent and Recoverable

- Database and Transaction logs persisted on local disk or flash storage
- Replication to standby and DR systems

Extremely Fast

- Microseconds response time
- Very high throughput

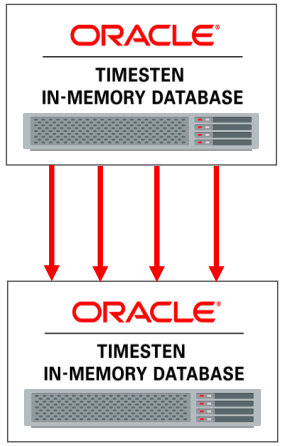

Highly Available

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

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

Most Widely Used Relational In-Memory Database

Deployed by Thousands of Companies

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

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

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

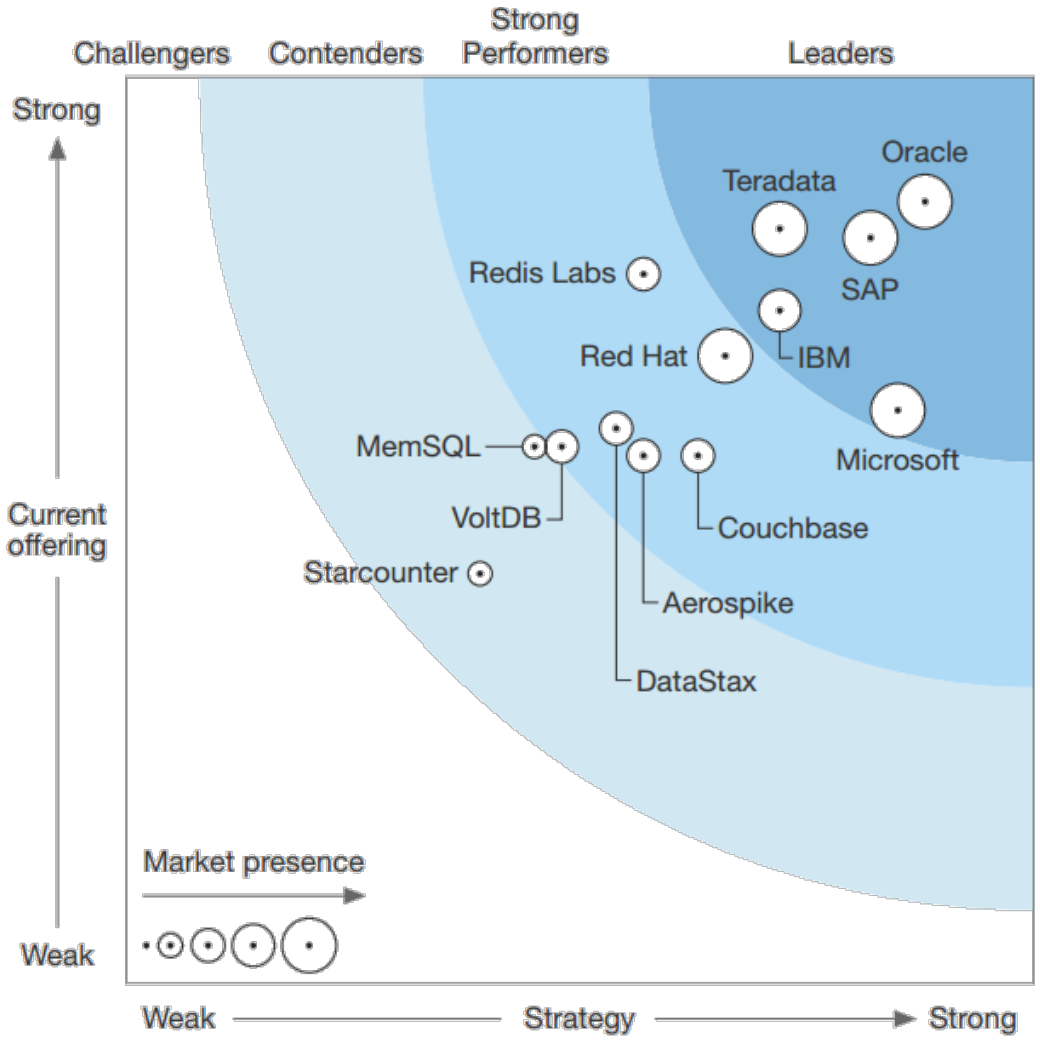

http://www.oracle.com/us/corporate/analystreports/forrester-imdb-wave-2017-3616348.pdf

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.

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

Single Database Image

- Database size not limited by memory
- Table data distributed across all elements
  - All elements are equal
- Connect to **any** element and access **all** data
  - Distributed queries, joins & transactions
- No need to de-normalize data model

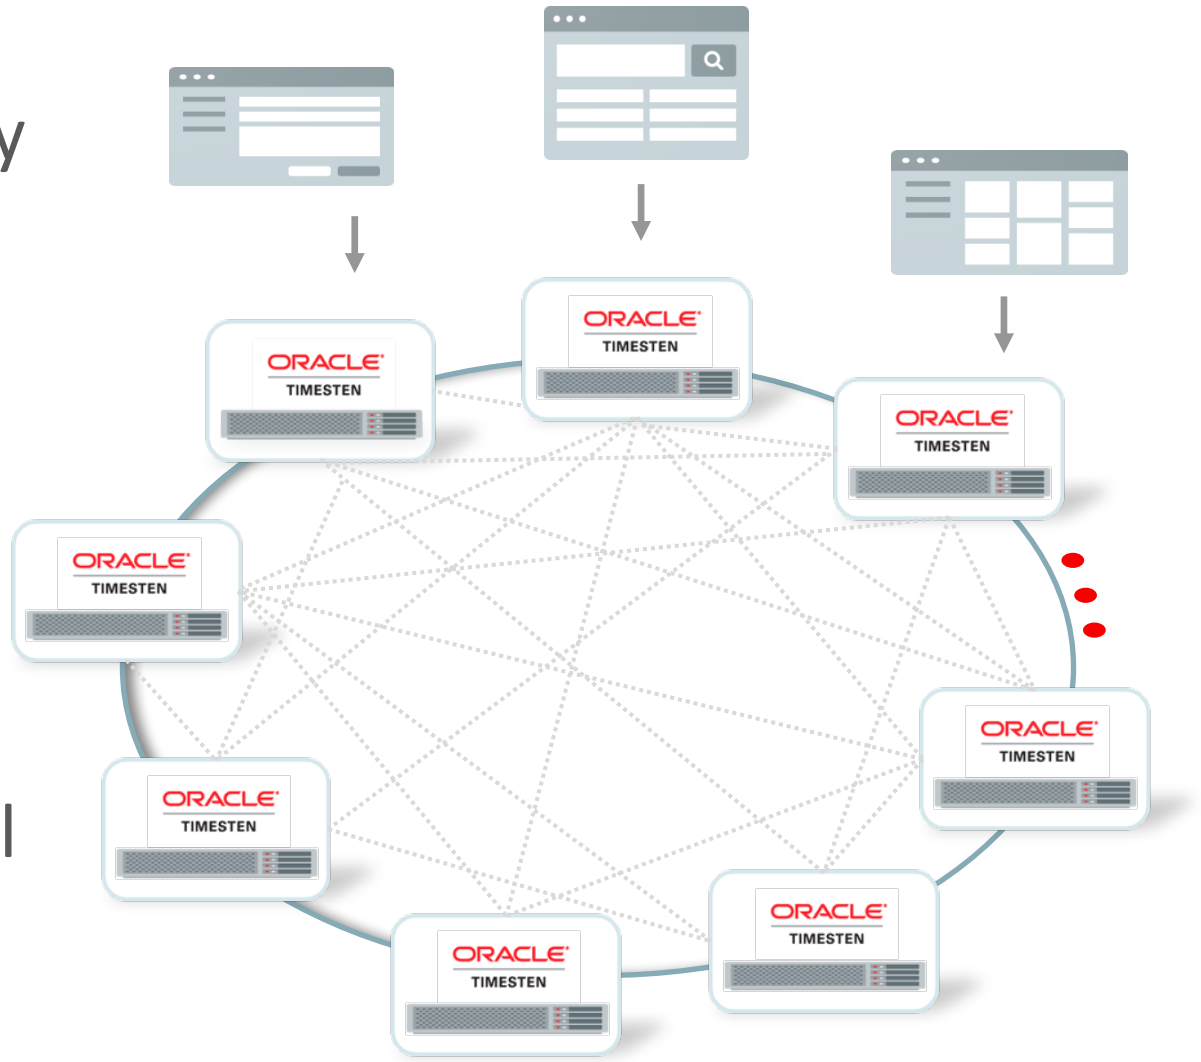

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. |

High Availability and Maximum Throughput

K-Safety, All Active

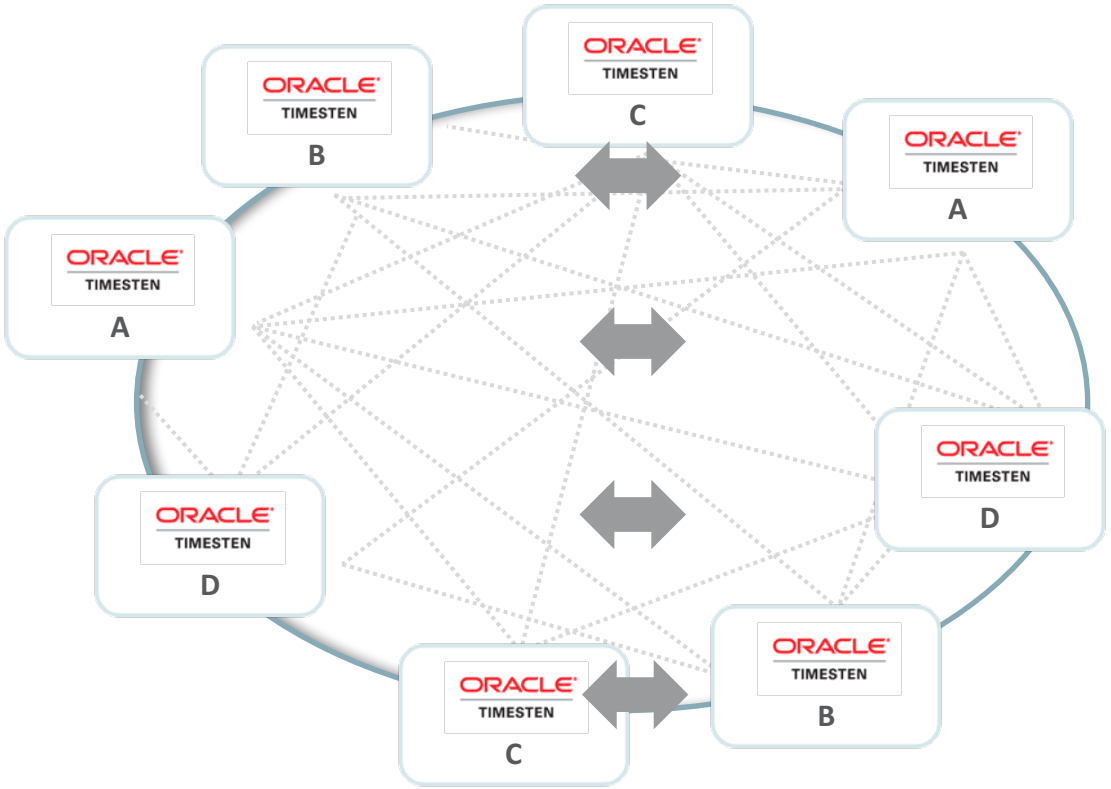

- Built-in HA via multiple copies of the data (K-safety)
  - Automatically kept in sync
- **All** replicas are **active** for **reads** and **writes**
  - Double the compute capacity
- Transactions can be initiated from and executed on any replica

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

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

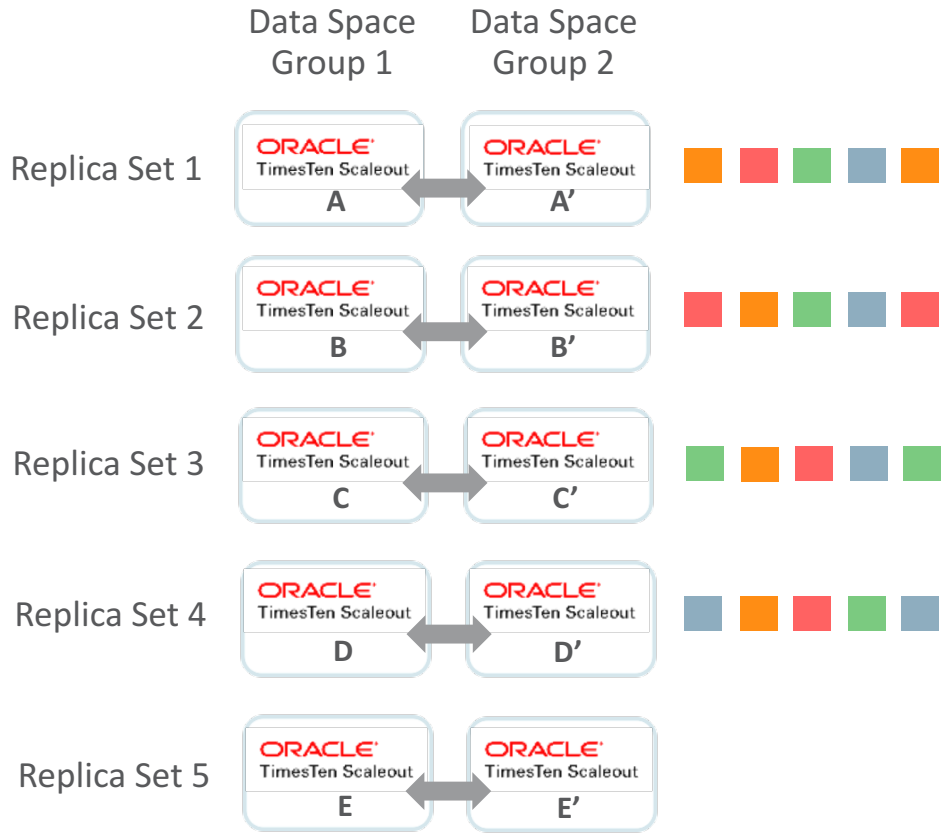

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. |

Database Fault Tolerance – No Application Down Time

Provided one entire copy of the database is available

- If multiple elements fail, applications will continue provided there is one complete copy of the database
- Recovery after failure is automatic
- If an entire replica set is down, application can **explicitly** choose to accept partial results

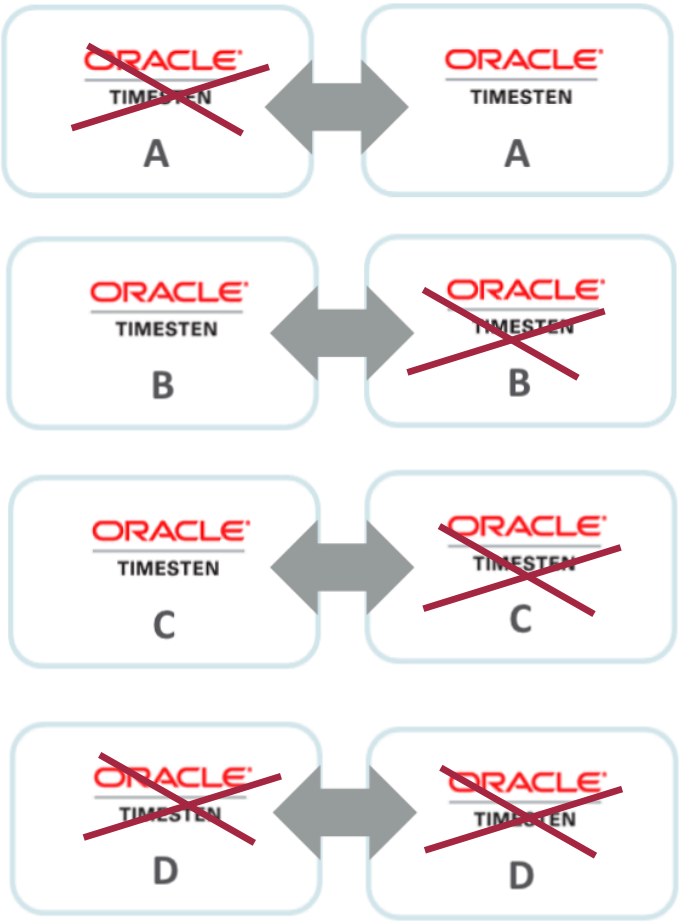

Copyright © 2018, Oracle and/or its affiliates. All rights reserved. |

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