
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.

The Next Generation of In-Memory Database

Oracle TimesTen Velocity Scale In-Memory Database

September 18–22, 2016
San Francisco

Doug Hood
TimesTen Product Manager

November 17, 2016

ORACLE

Agenda

1. Introduction to TimesTen
2. TimesTen Velocity Scale In-Memory Database
3. Demo

Oracle TimesTen – Pure In-Memory Relational Database

20 Years of Extreme Performance

Pre-Oracle acquisition

- 1998 First commercial In-Memory RDBMS
- HA Replication
- Online Upgrades
- Application-tier Cache for Oracle Database

TimesTen 6
TimesTen 7

- Oracle RAC integration
- National Language Support
- Oracle Data Types support
- SQL Developer Integration
- Enterprise Manager integration

TimesTen 11*g*
11.2.1

- PL/SQL and OCI Support
- Oracle Clusterware Integration
- Cache Grid for Scale Out
- ODP .NET Support
- BLOB, CLOB, NCLOB data types

TimesTen 11*g*
11.2.2

- Parallel Replication
- In-Memory Analytics
- Columnar Compression
- Index Advisor
- Oracle R Support
- In-Memory Star Join
- Oracle Golden Gate Integration

TimesTen 11.2.2.x
Enhancements

- Parallel data import from Oracle Database
- Parallel database restart
- Highly concurrent range indexes
- Parallel Replication with commit order optimization

Oracle TimesTen In-Memory Database

Relational Database

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

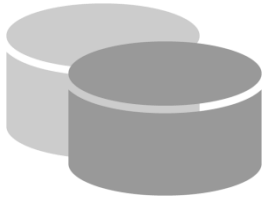

Persistent and Recoverable

- 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

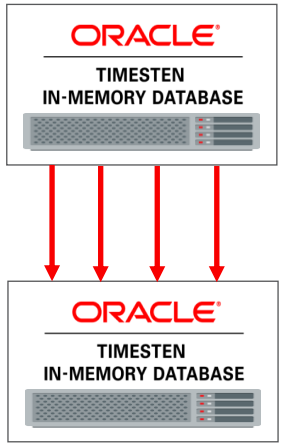

Highly Available

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

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)

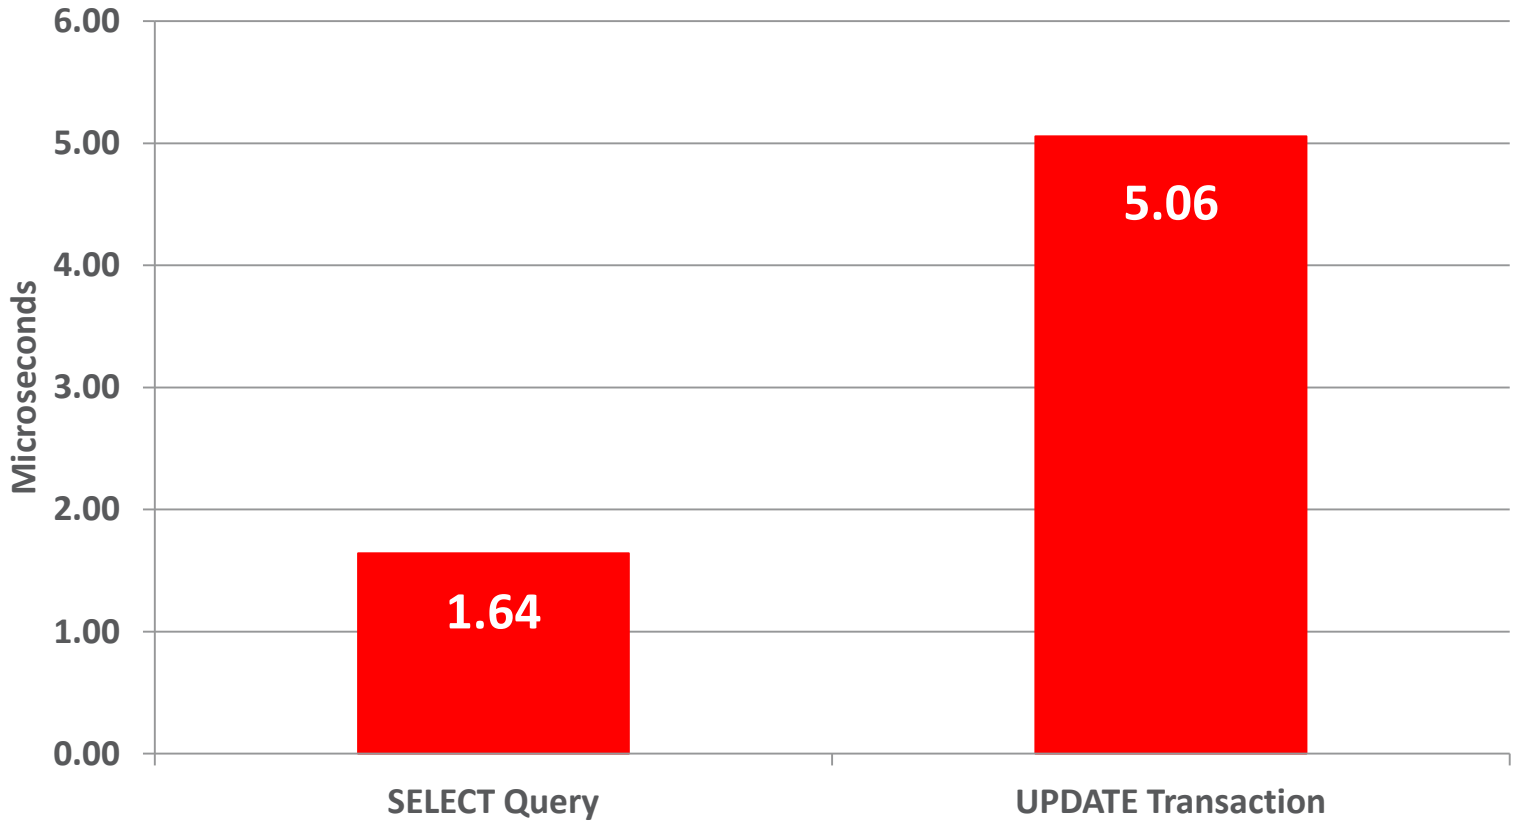

Performance - Throughput

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)

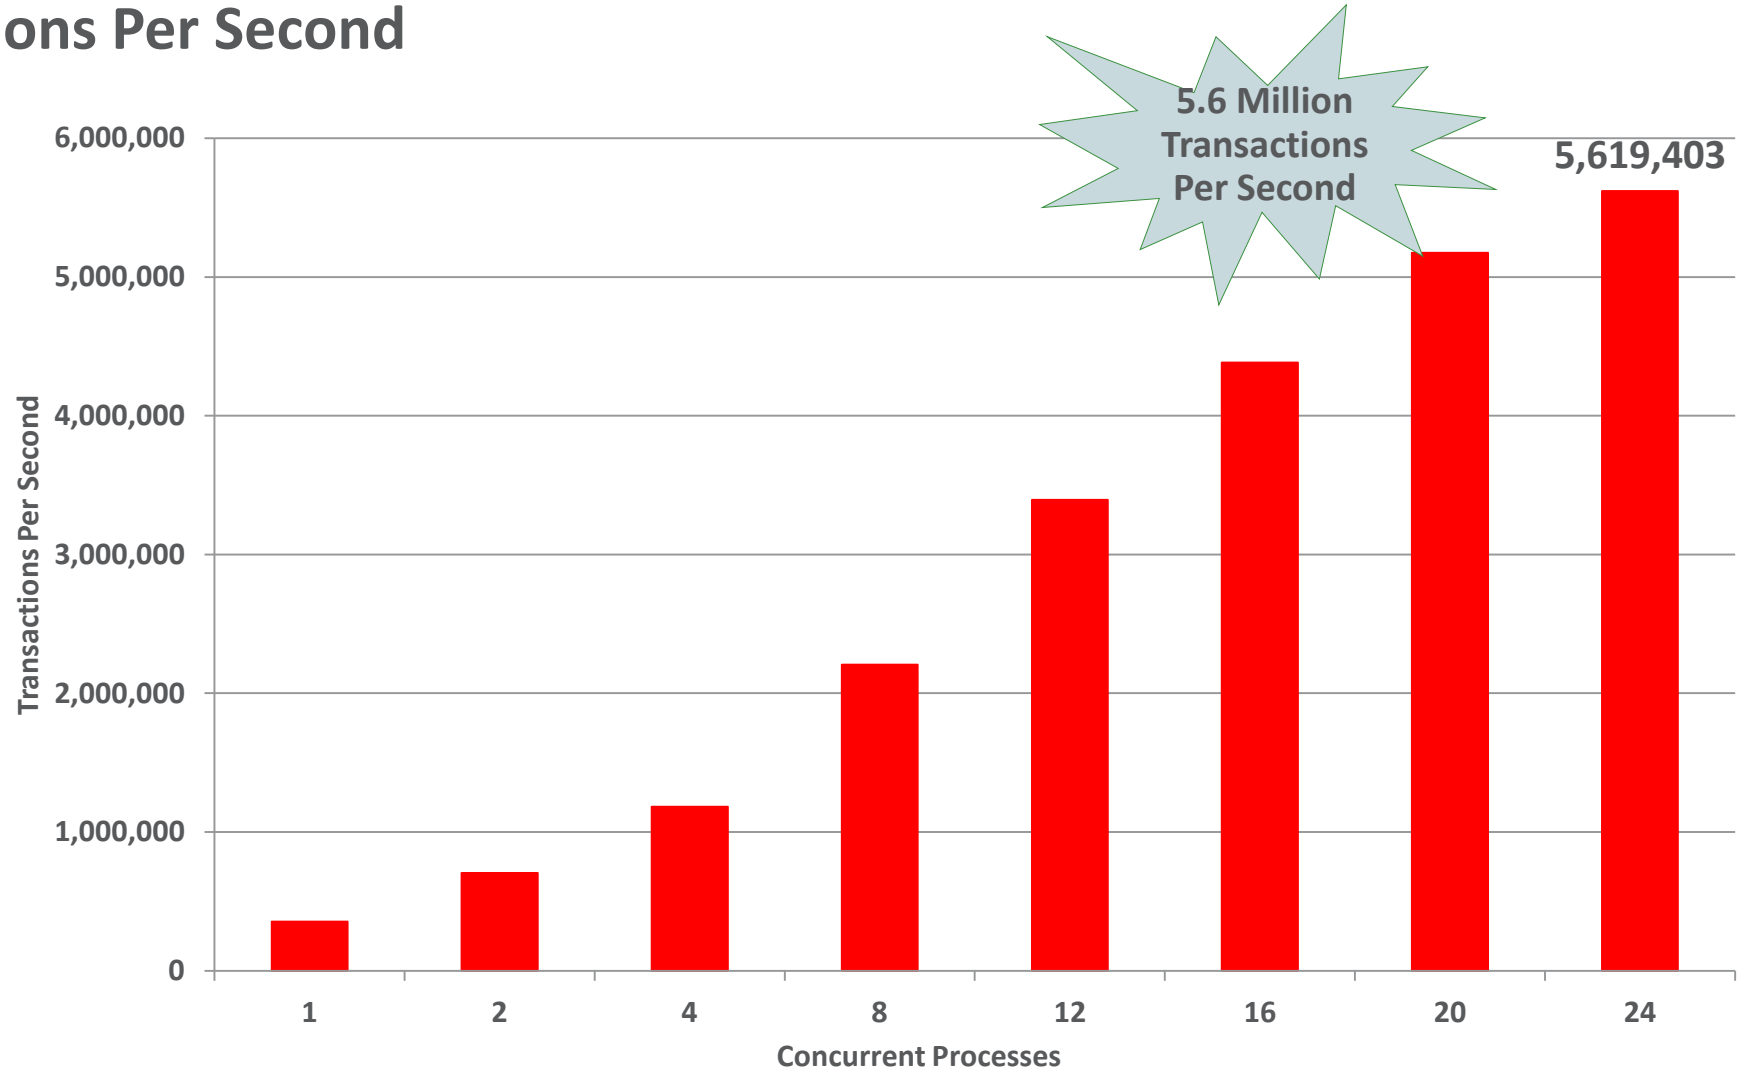

Most Widely Used Relational In-Memory Database

Deployed by Thousands of Companies

1 Introduction to TimesTen

2 TimesTen Velocity Scale In-Memory Database

3 Demo

Oracle TimesTen Velocity Scale In-Memory Database

- **Scale-out**, shared nothing architecture
- Data location **transparency**
- Built-in **high availability** via K-safety
- **Elastic** scalability
- **Easy** to **deploy** and **manage**
- **Easy** application **development** for **OLTP** workloads
- Supports **analytic** workloads
- **On-premises** or **Cloud** deployment

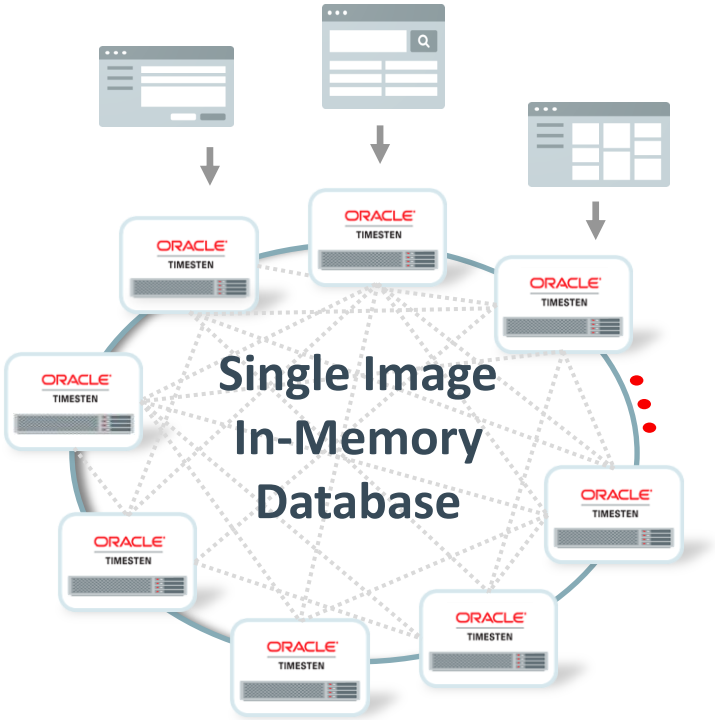

Database Element
Unit of Persistence and Recovery

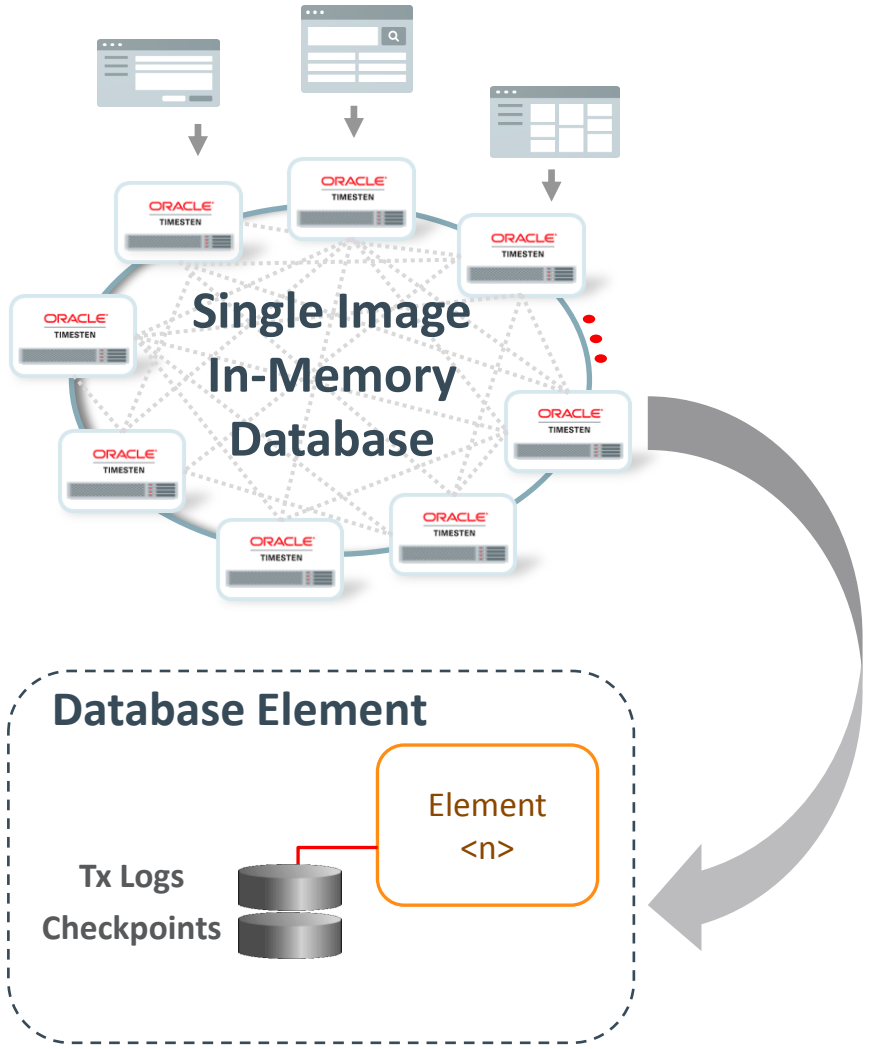

- Each Element has its own set of checkpoint files and transaction log files for persistence
- The Element is the smallest unit for database persistence, failure recovery and high availability

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

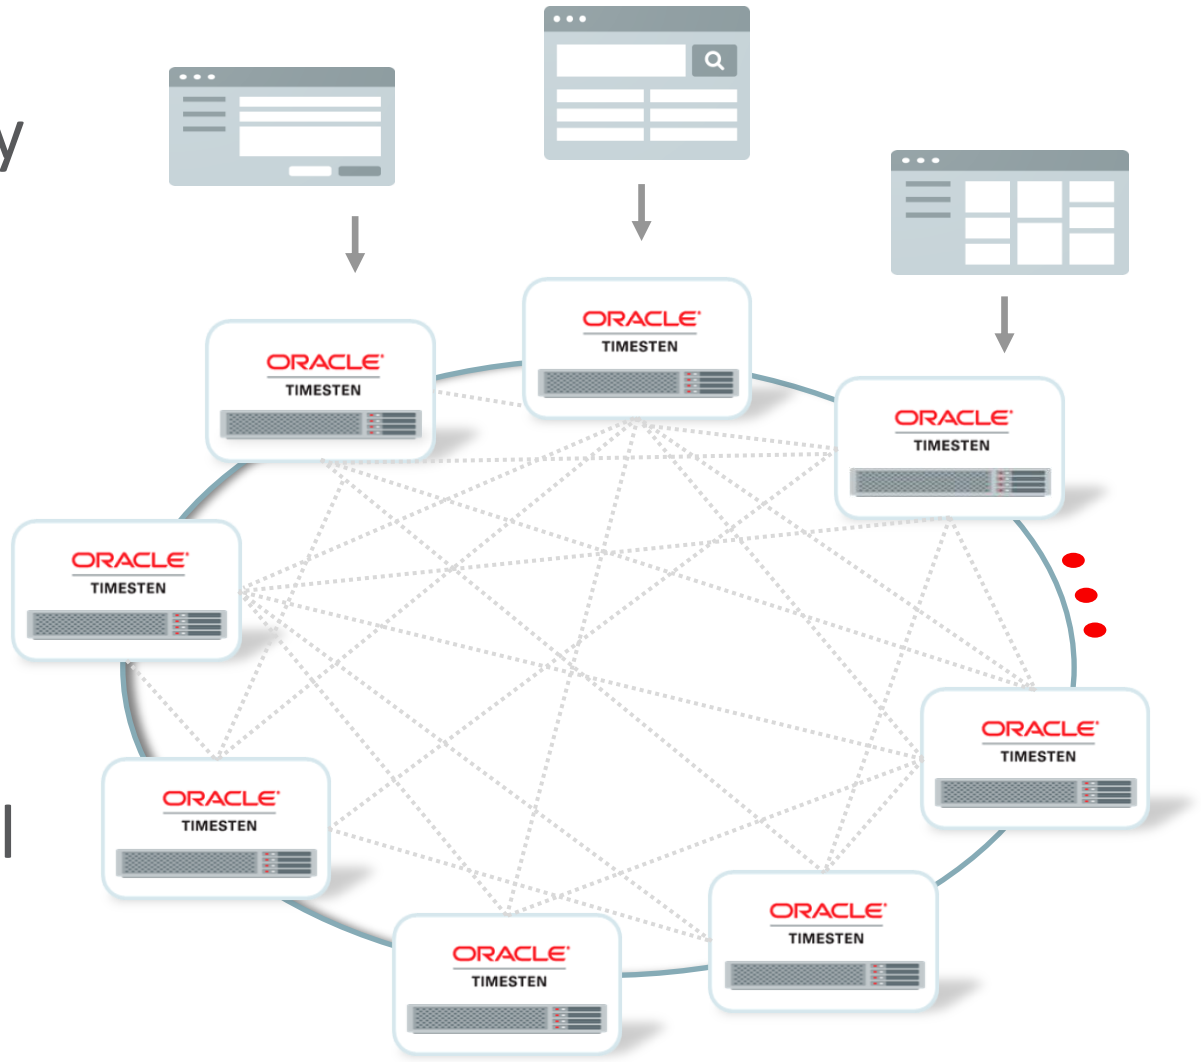

Data Distribution Methods

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

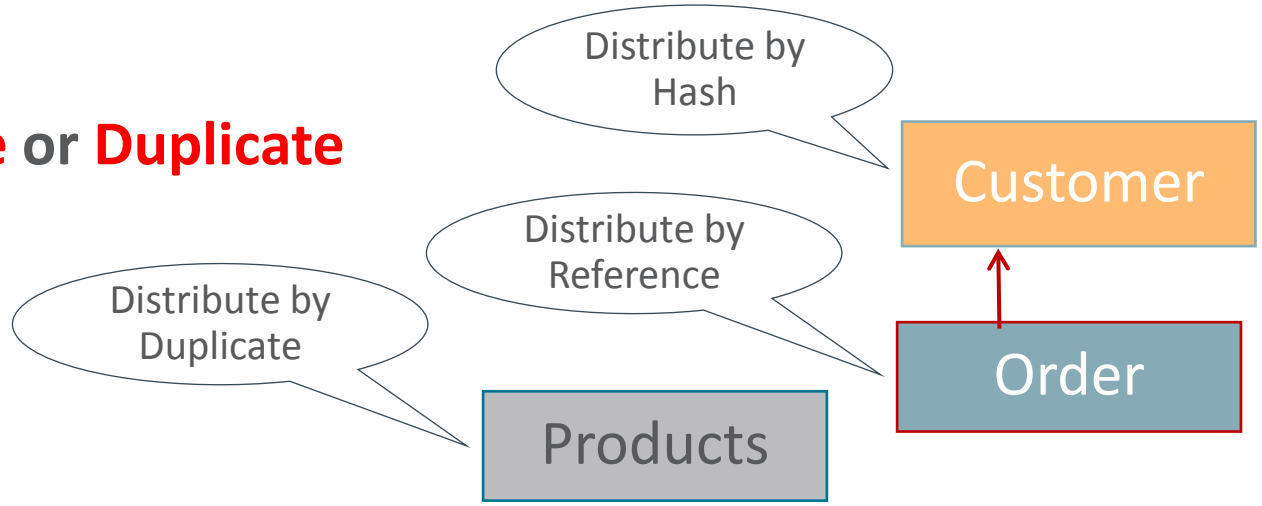

Element 1

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

High Availability and Maximum Throughput

K-Safety, All Active

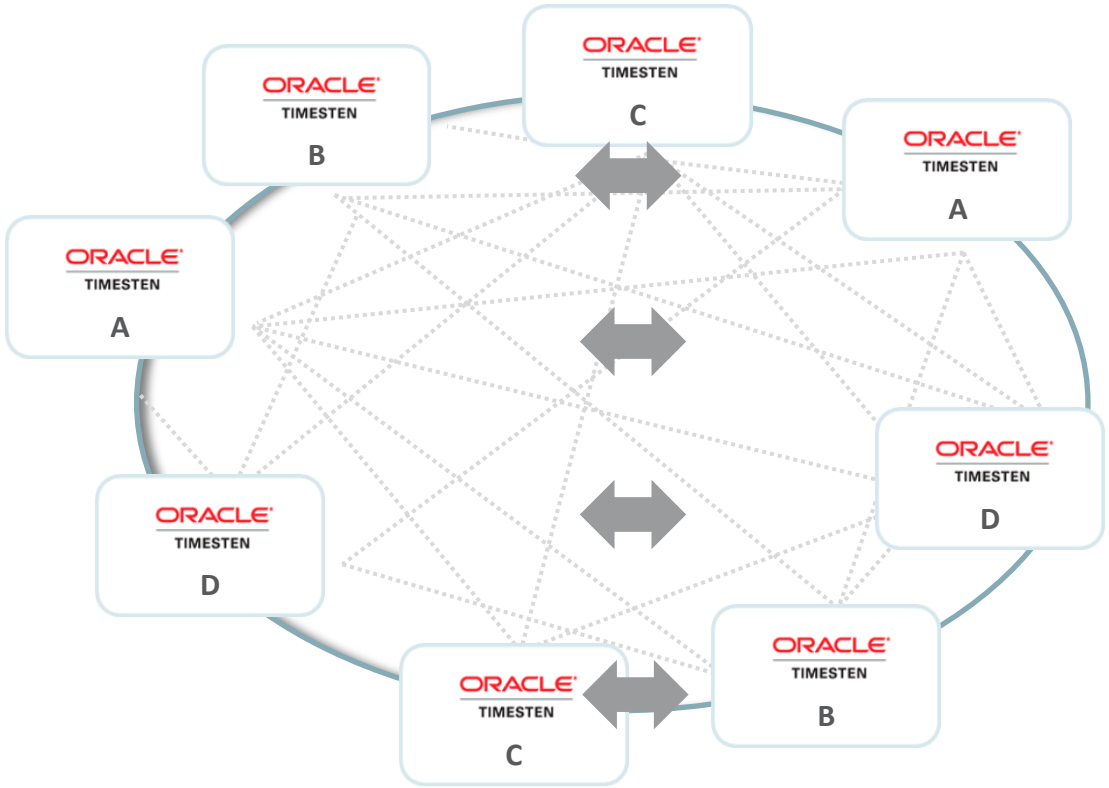

- Built-in HA via multiple copies of the data (K-safety)
  - Automatically kept in sync
- **All** replicas are **active** for **reads** and **writes**
  - Double the compute capacity
- Transactions can be initiated from and executed on any replica

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

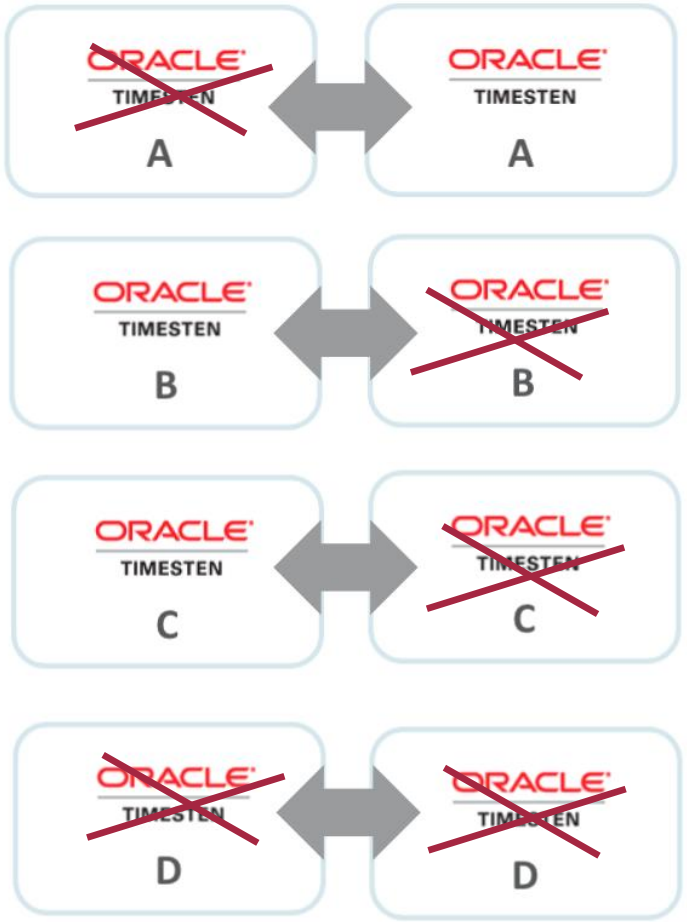

Elastic Scalability

Expand and shrink the database based on business needs

Adding and removing DB elements

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

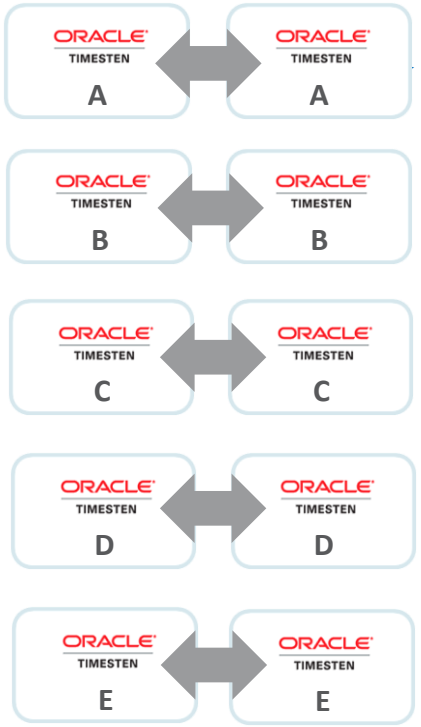

Deployment and Management

Single point of Control, via command line or SQL Developer

- Single point of configuration and deployment
  - Management instance(s)
- Powerful command line tools
  - ttGridAdmin and ttGridRollout
- Extensive support in SQL Developer as well
  - Velocity Scale grid rollout, monitoring and management
  - SQL and table editor support
  - Data loading

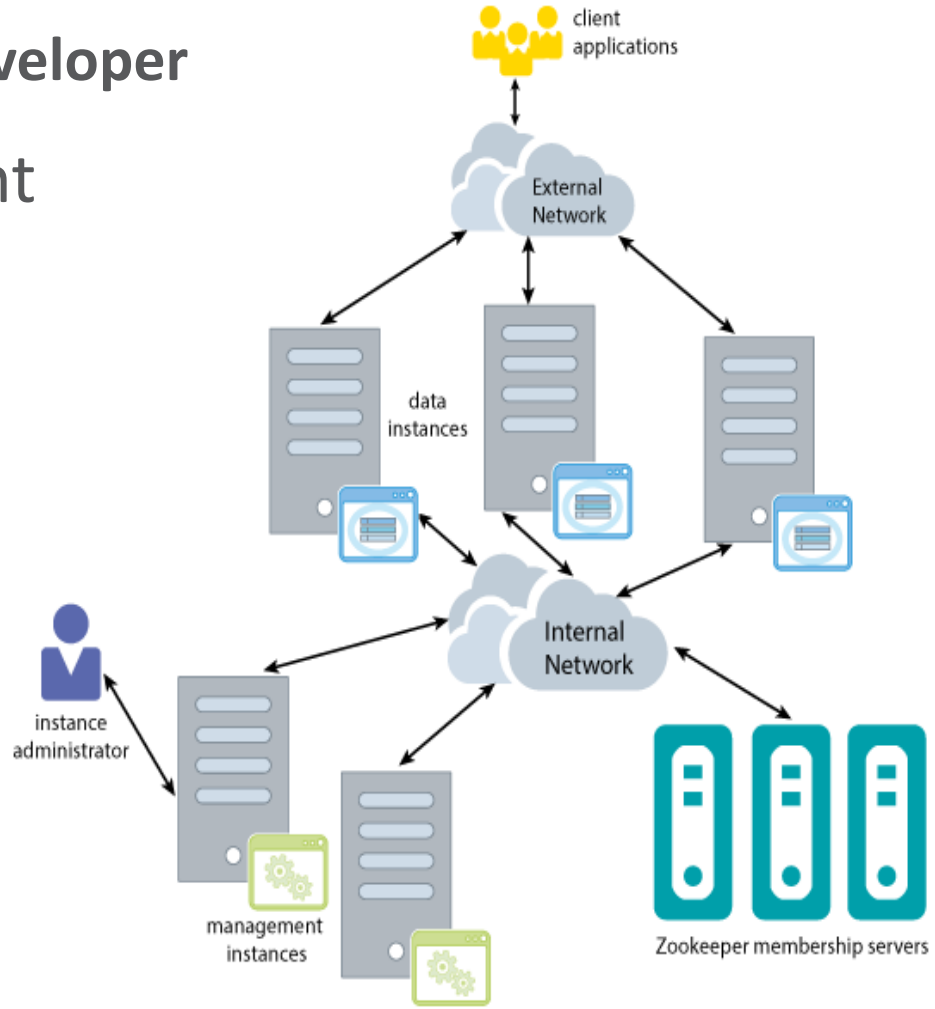

Common Database Performance Myths

- **Myth 1:** You need a NoSQL DB to scale
- **Myth 2:** You need to abandon ACID txns to scale
- **Myth 3:** You need to de-normalize [no joins] to scale
- **Myth 4:** You need to manually shard your DB to scale
- **Myth 5:** You need to re-write your apps to scale

- Velocity Scale is proof that these myths are false
- Velocity Scale offers performance without compromise

TimesTen Velocity Scale IMDB vs Cassandra

YCSB (Yahoo Cloud Service Benchmark) – 50% reads 50% writes [Workload A]

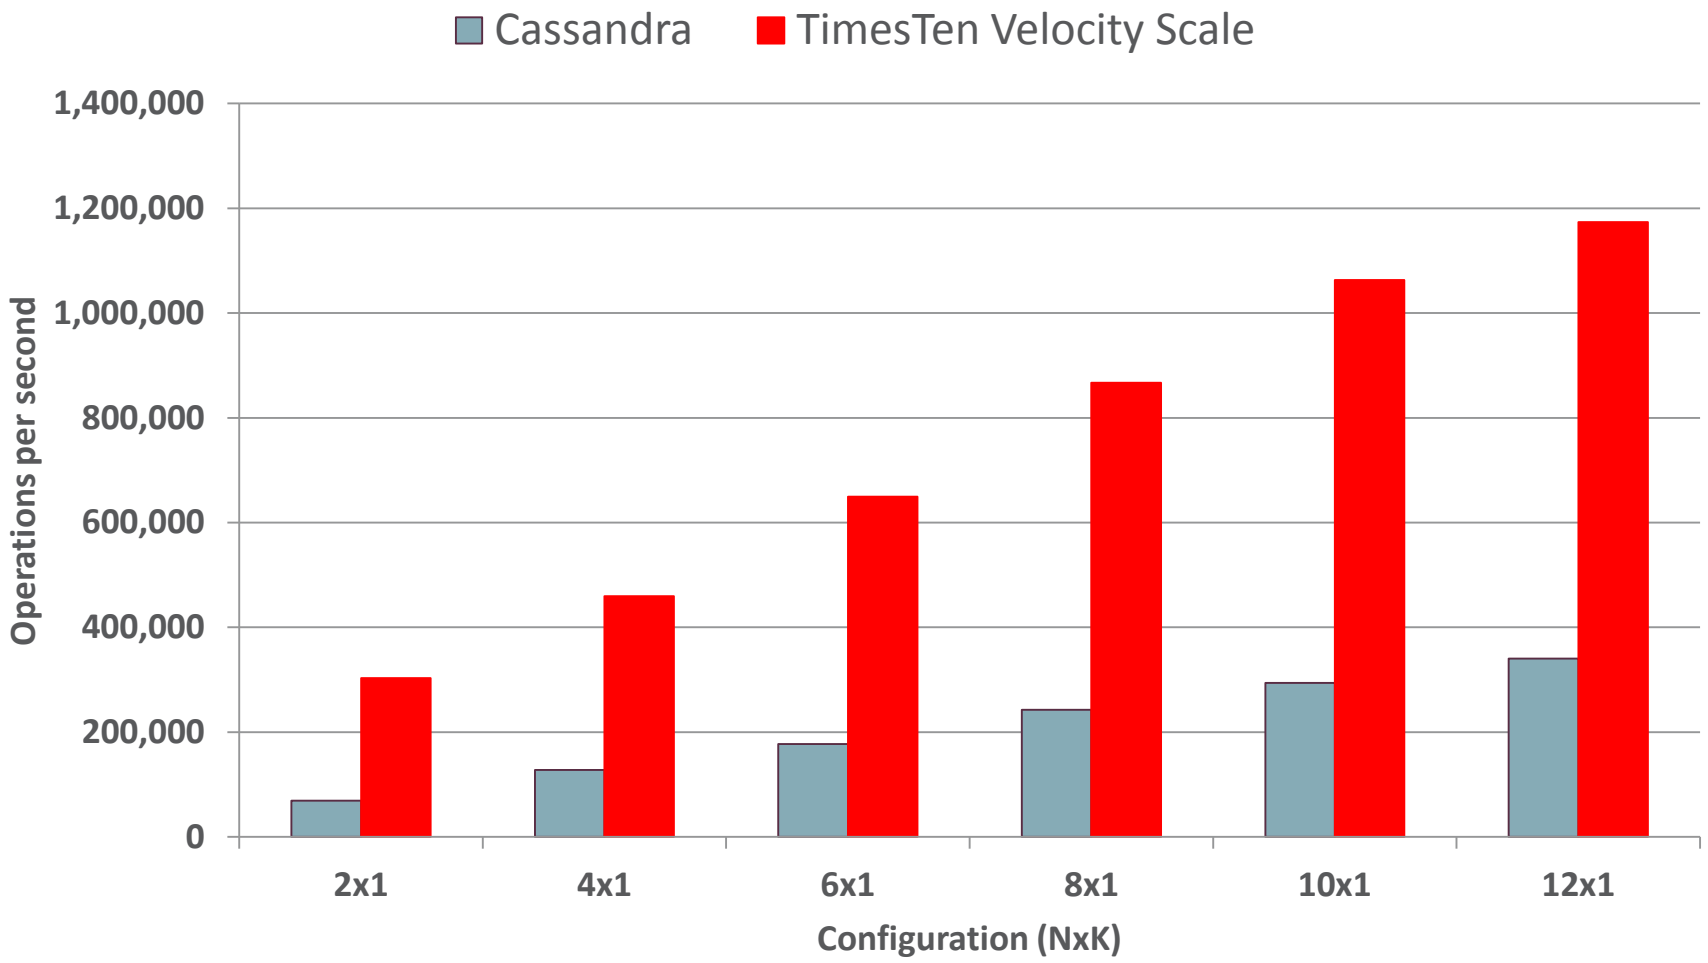

YCSB N*10M Rows
Servers: 12 x X5-2
Clients: 2 x ODA
Network: 10Gbps

TimesTen Velocity Scale IMDB vs Cassandra

YCSB (Yahoo Cloud Service Benchmark) – 95% reads 5% writes [Workload B]

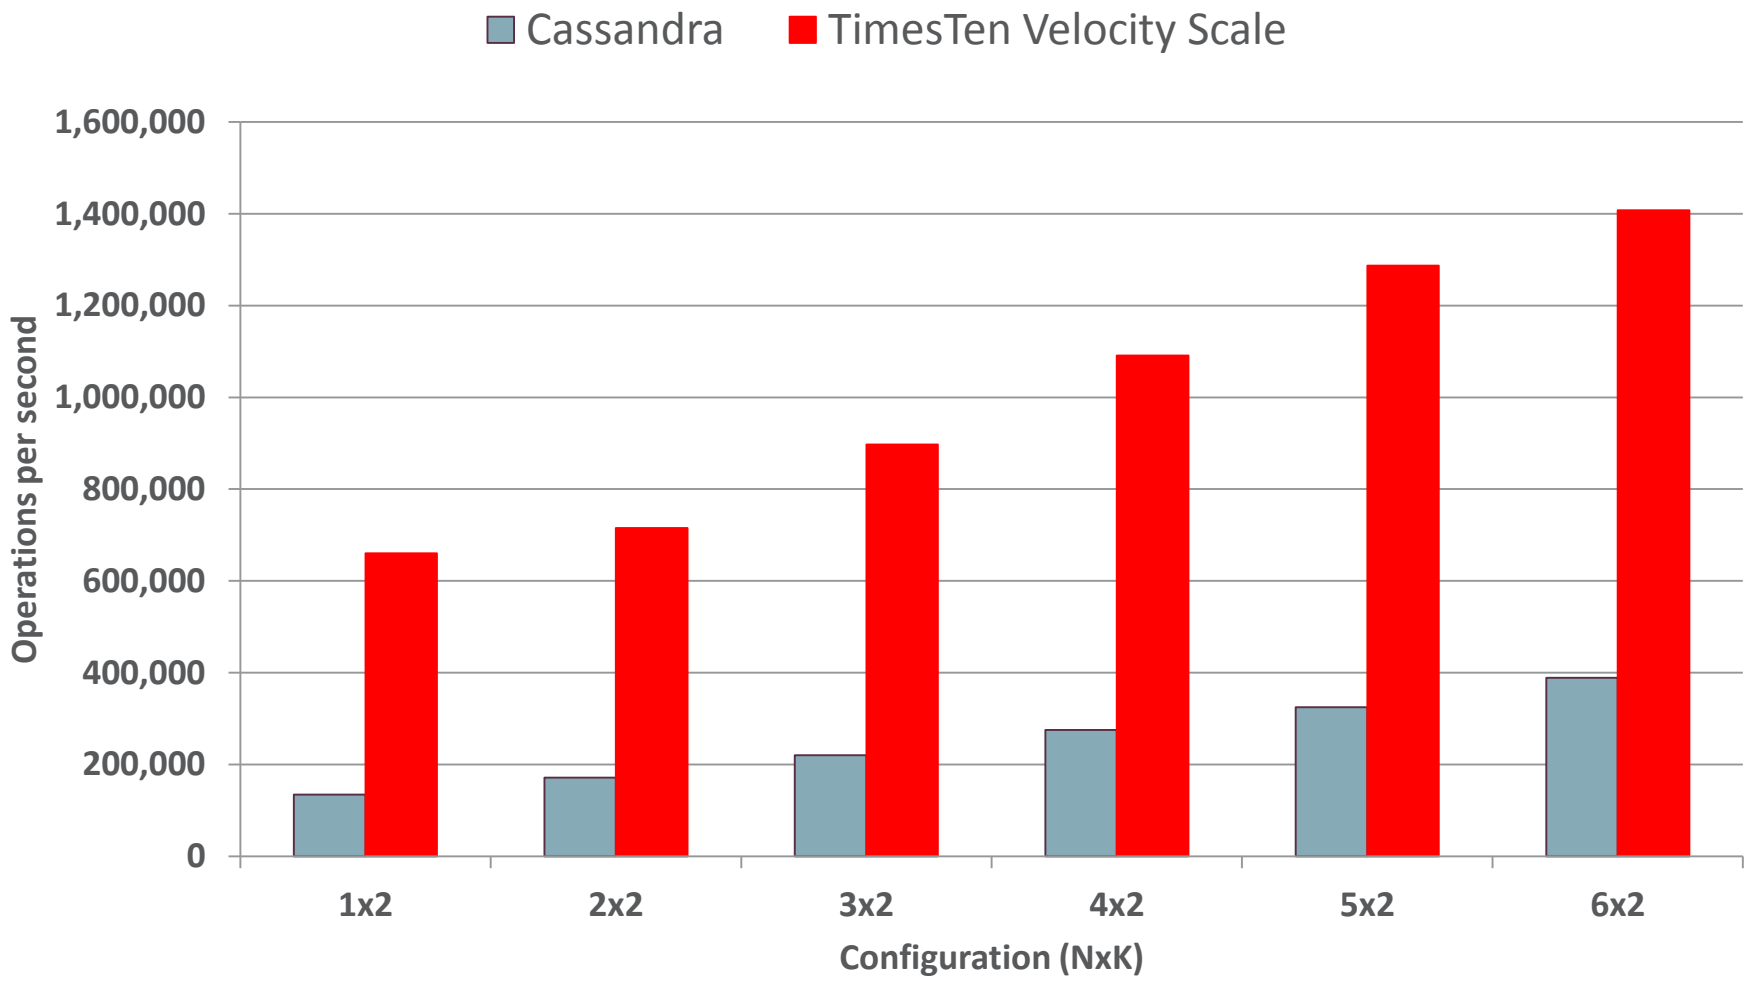

YCSB N*10M Rows
Servers: 12 x X5-2
Clients: 2 x ODA
Network: 10Gbps

TimesTen Velocity Scale IMDB vs Cassandra

YCSB (Yahoo Cloud Service Benchmark) – 95% reads 5% writes [Workload B]

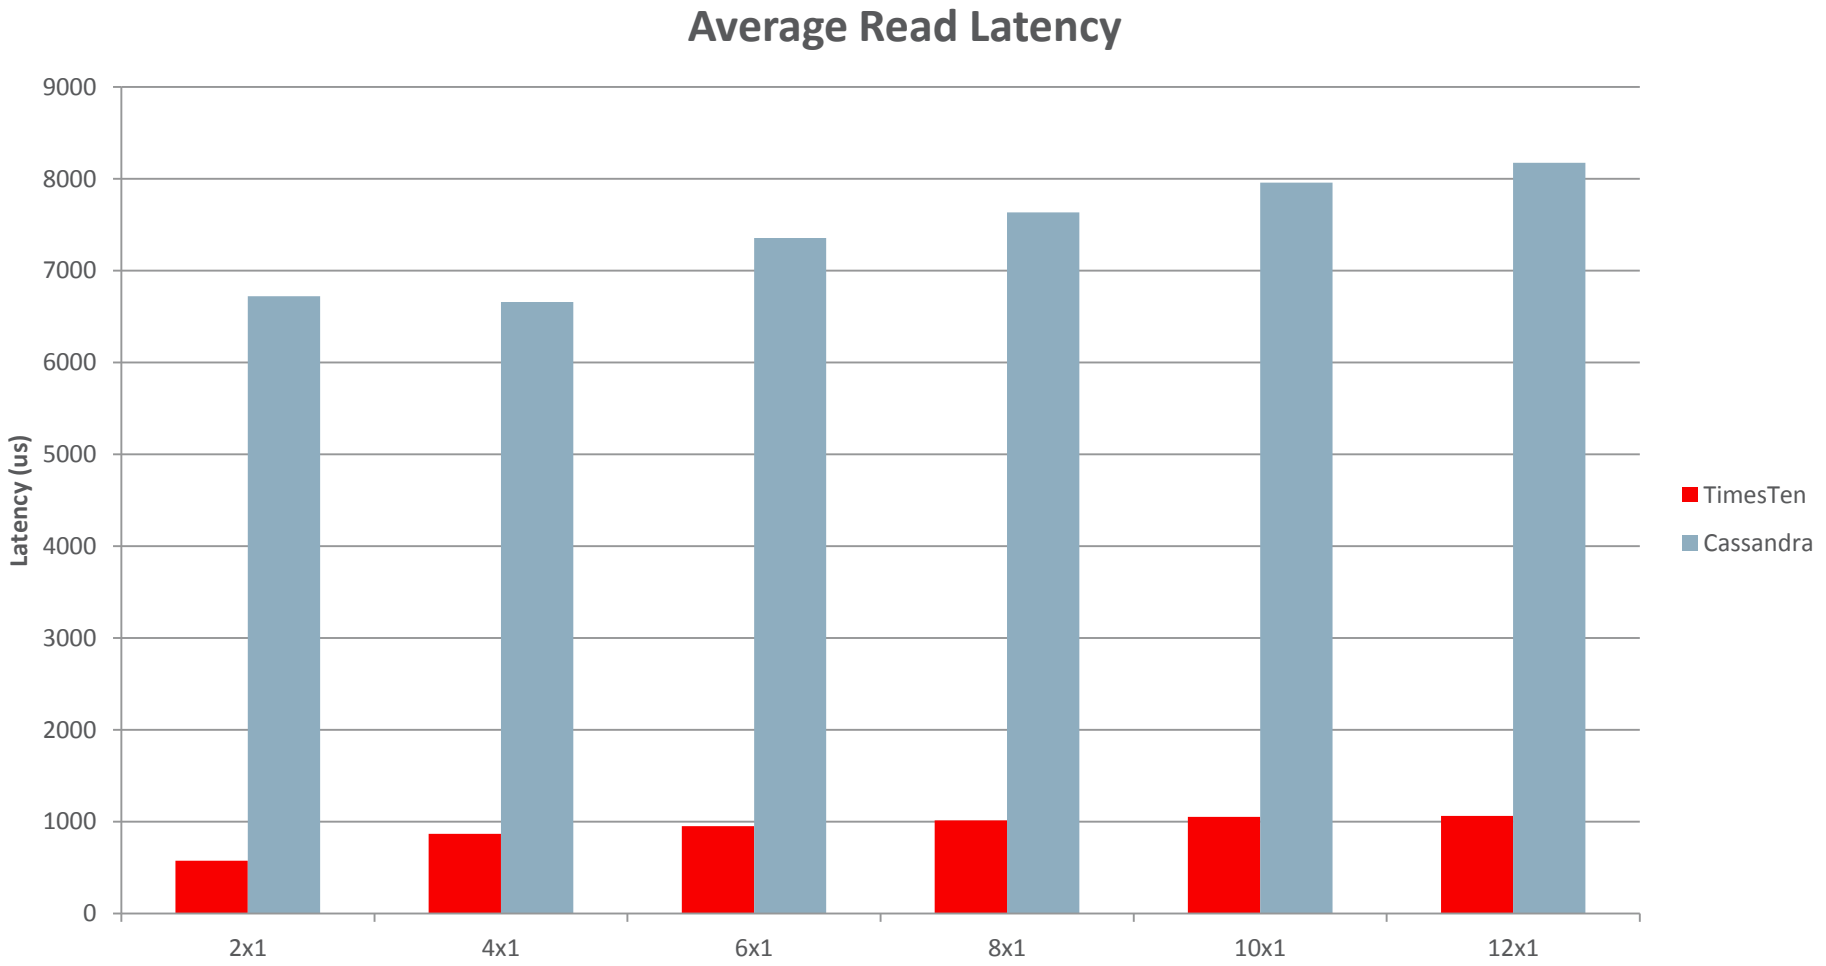

YCSB N*10M Rows
Servers: 12 x X5-2
Clients: 2 x ODA
Network: 10Gbps

VelociScale DB – Recalc priority Flow

Complex Update Txn (9 table join, hundreds of millions of rows)

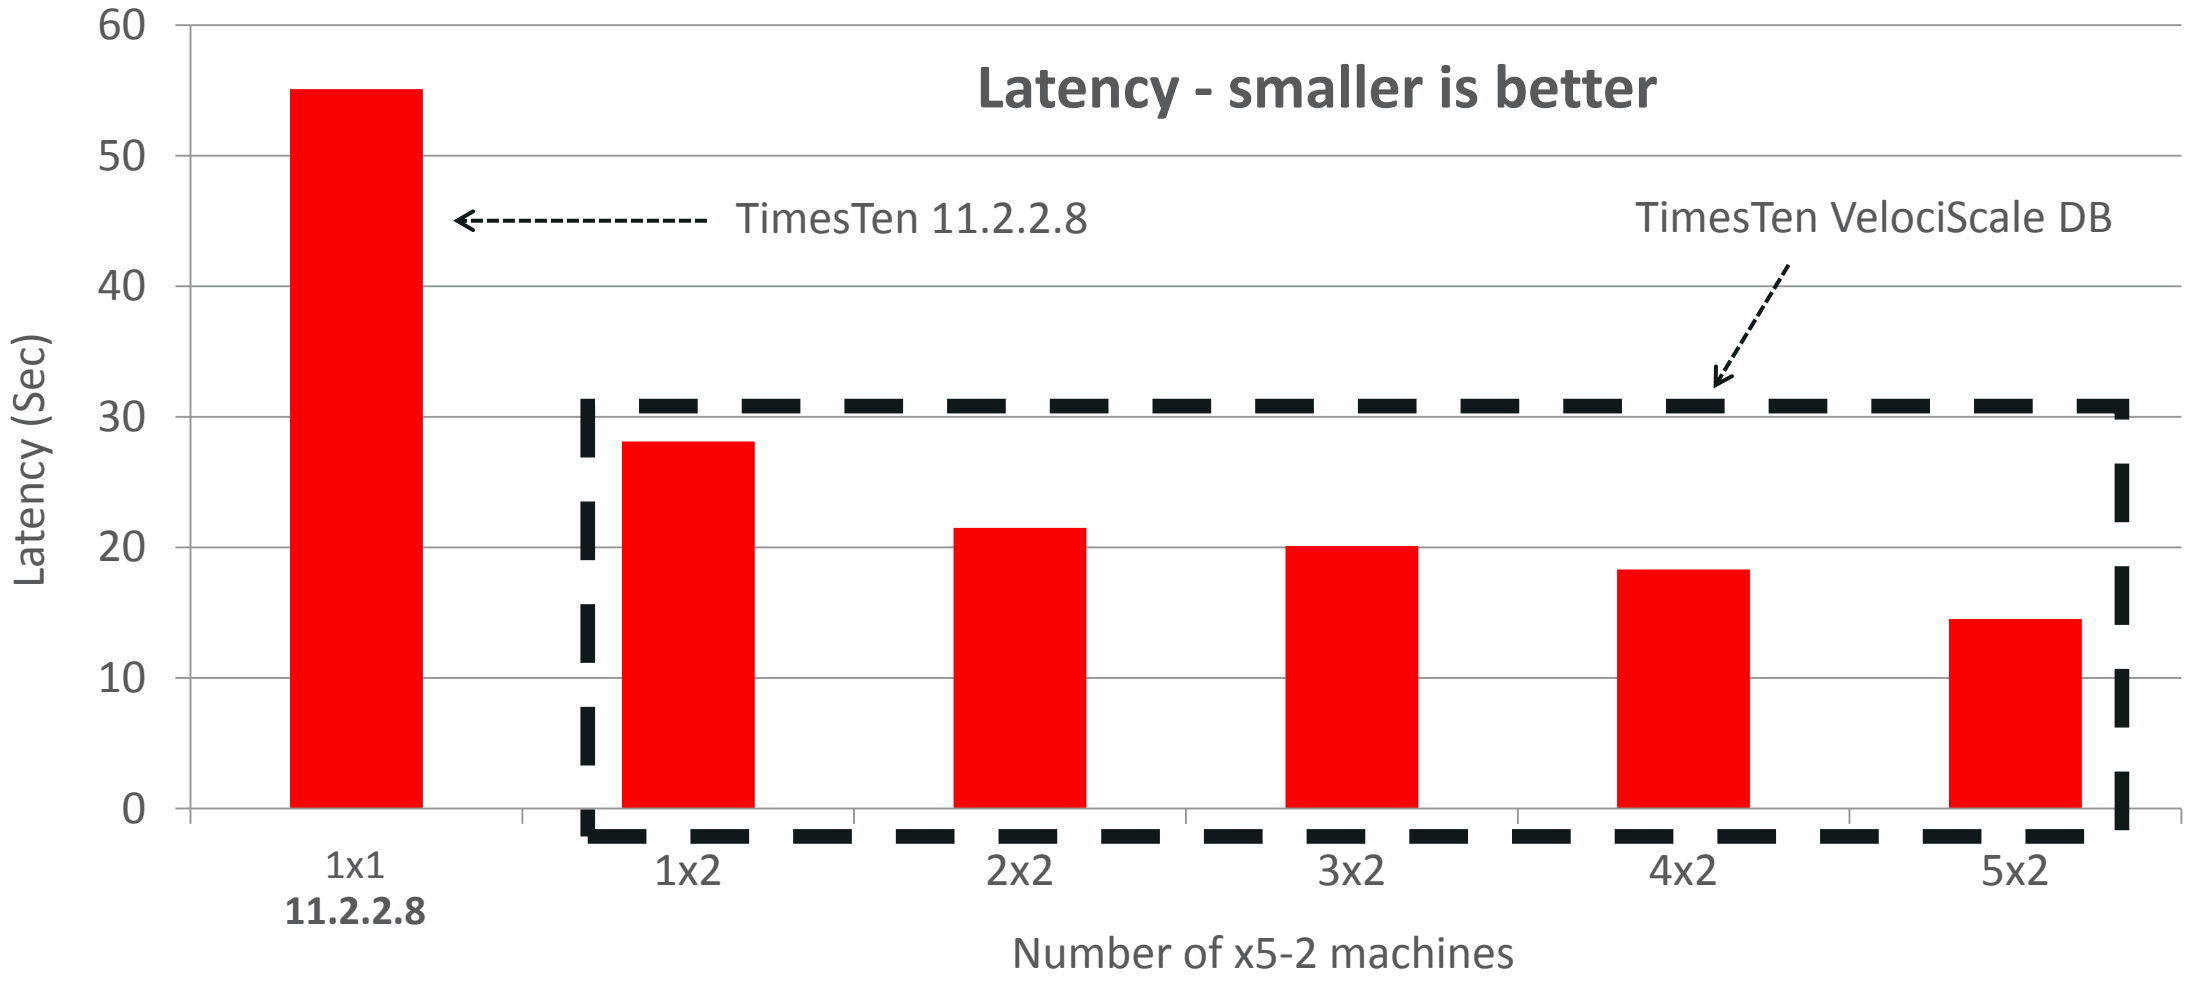

Application Development

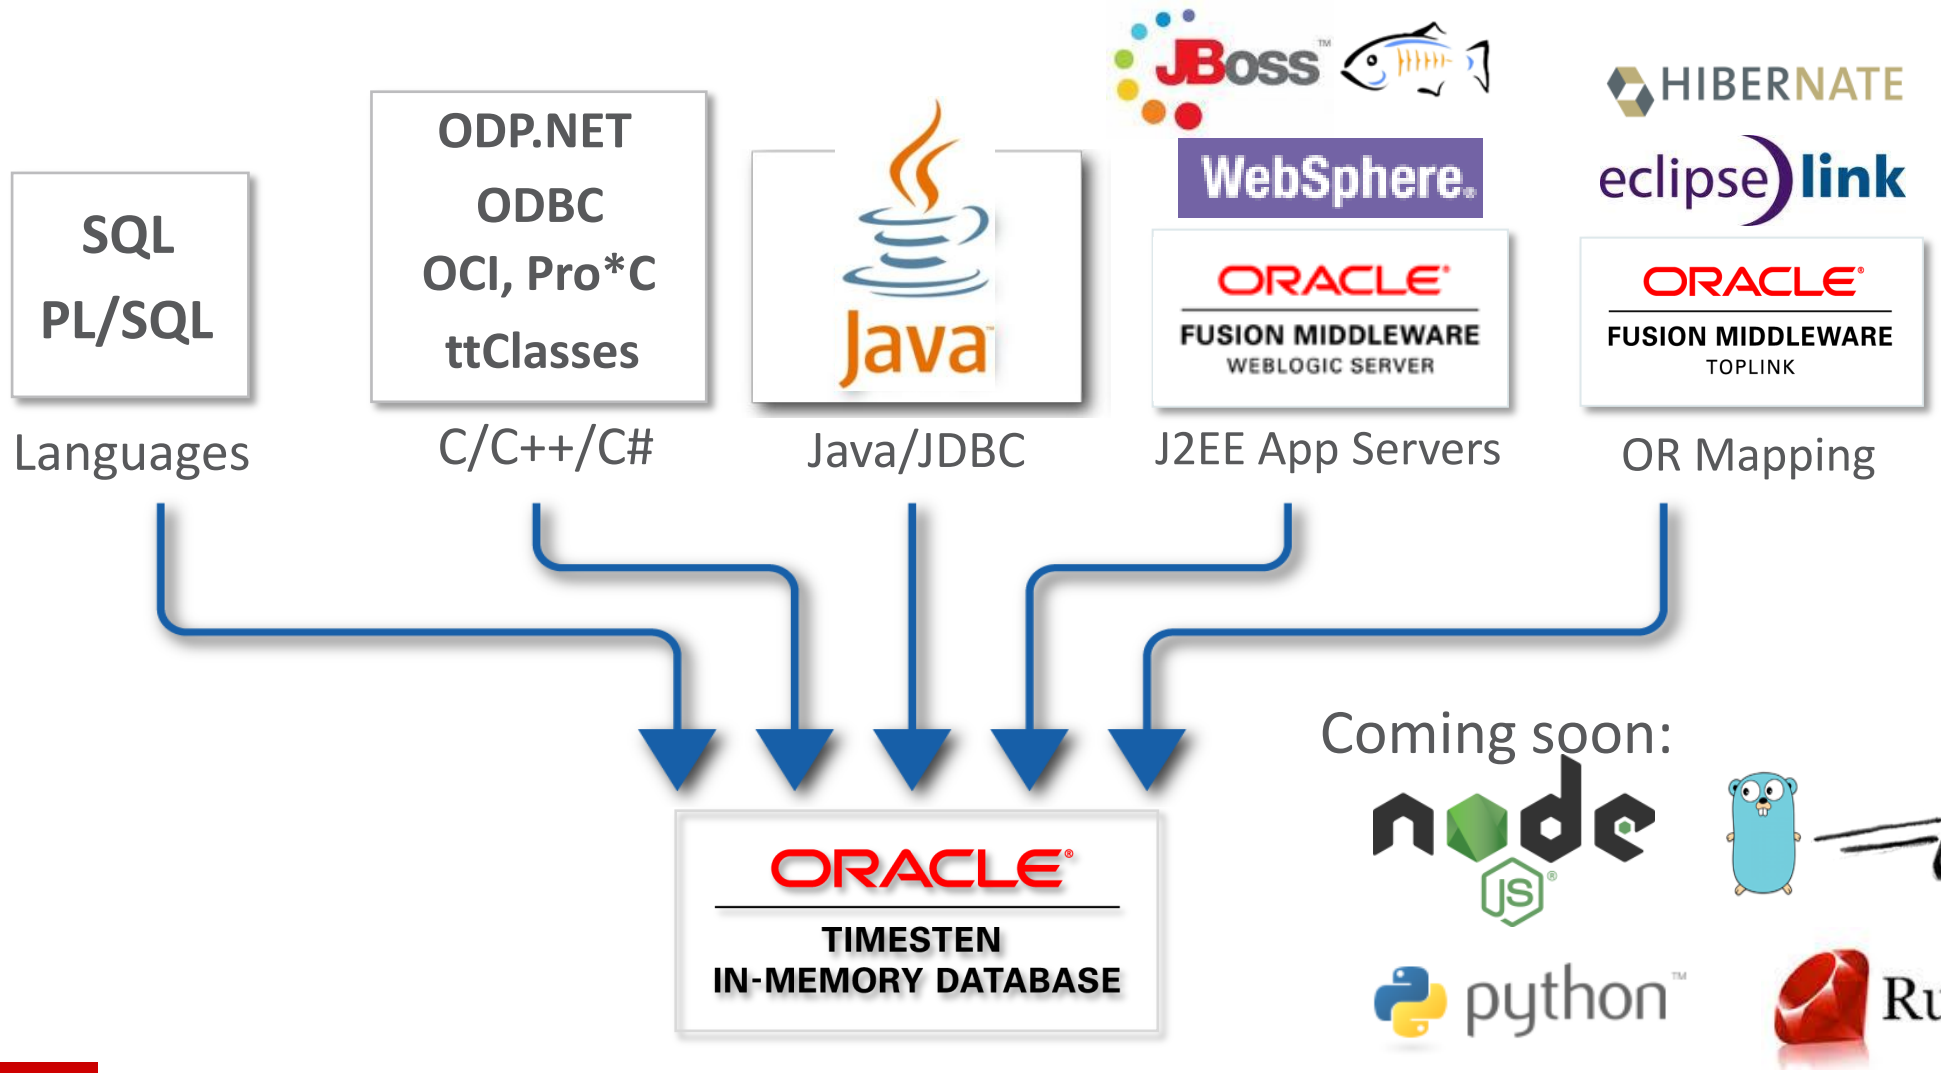

Developing Applications for Velocity Scale

- **Existing** RDBMS and TimesTen skills apply
  - Same SQL, same APIs, same languages
- **No application code changes** needed for existing TimesTen applications
- Need to consider table **distribution types**
  - Hash may not be optimal for all cases
- Use **Index Advisor** and **MVs** if SQL joins are the bottleneck

1 Introduction to TimesTen

2 TimesTen Velocity Scale In-Memory Database

3 Demo

Oracle TimesTen Velocity Scale In-Memory Database

- Scale-out, shared nothing architecture with single database image
- Data location transparency
- Built-in high availability via K-safety
- Elastic scale-out
- Easy to deploy and manage
- Easy application development for OLTP and Analytic workloads
- On-premises or Cloud deployment
- *You **can** have all of **elastic scalability**, **HA**, **high performance**, **ACID** transactions, **full consistency** and **full featured SQL***

Integrated Cloud

Applications & Platform Services


ORACLE®