

ORACLE®

Oracle TimesTen In-Memory Database and TimesTen Velocity Scale

Sam Drake
Architect, Oracle TimesTen In-Memory Database

NYOUG – 14 Sep 2016

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.

AGENDA

1. TimesTen Introduction
2. TimesTen Velocity Scale
3. Q&A

Oracle TimesTen In-Memory Database

Microseconds Response Time in Application Tier

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 disk / flash
- Replication to standby and DR systems



Extremely Fast

- Microseconds response time
- Very high throughput



Compatible with Oracle Database

- Data types, PL/SQL, JDBC, OCI, ODP.NET, PHP, R
- Integrated with RAC, Data Guard, Enterprise Manager, SQL Developer, etc.

Oracle Database In-Memory vs Oracle TimesTen

- Database In-Memory Option
 - Feature of Oracle Database
 - Primarily intended for analytics workloads
 - Scans billions of rows/second
 - Optimized for set-oriented data processing

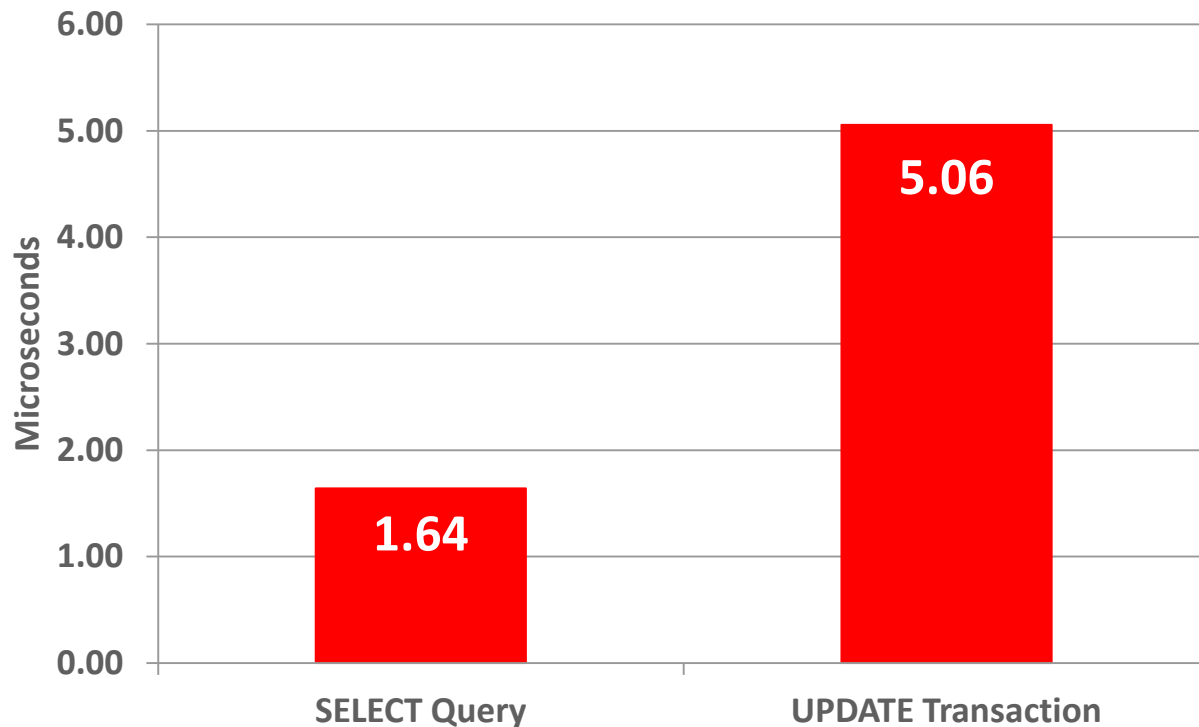


- TimesTen In-Memory Database
 - Runs in the application tier
 - Primarily for low-latency applications
 - Microsecond response time
 - Optimizes OLTP processing
 - E.g. insert a new stock trade, connect a cellular phone call



TimesTen In-Memory Database

Low Latency - Microseconds Response Time

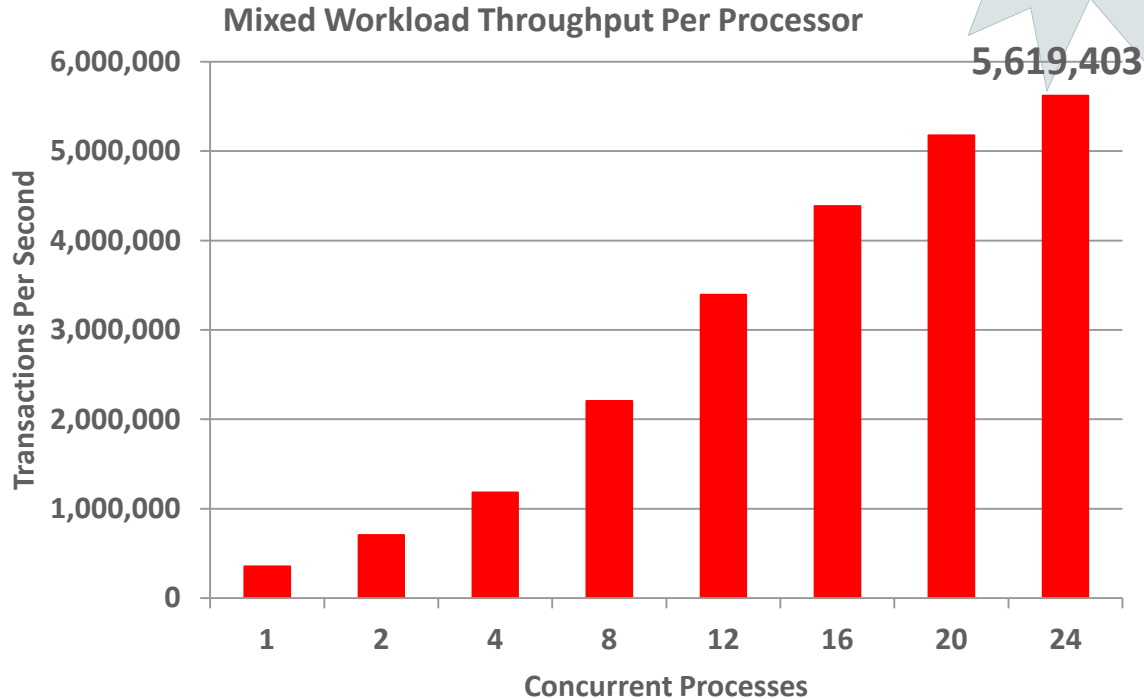


TPTBM Read and Update
E5-2699 v4 @ 2.20GHz
2 socket, 22 cores/socket,
2 threads/core
TimesTen 11.2.2.8.0
(100M rows, 17GB)

TimesTen In-Memory Database

5.6 Million Transactions Per Second Per Processor

**5.6 Million
Transactions
Per Second**



80-10-5-5 Workload = 80% select, 10% updates, 5% inserts, 5% deletes

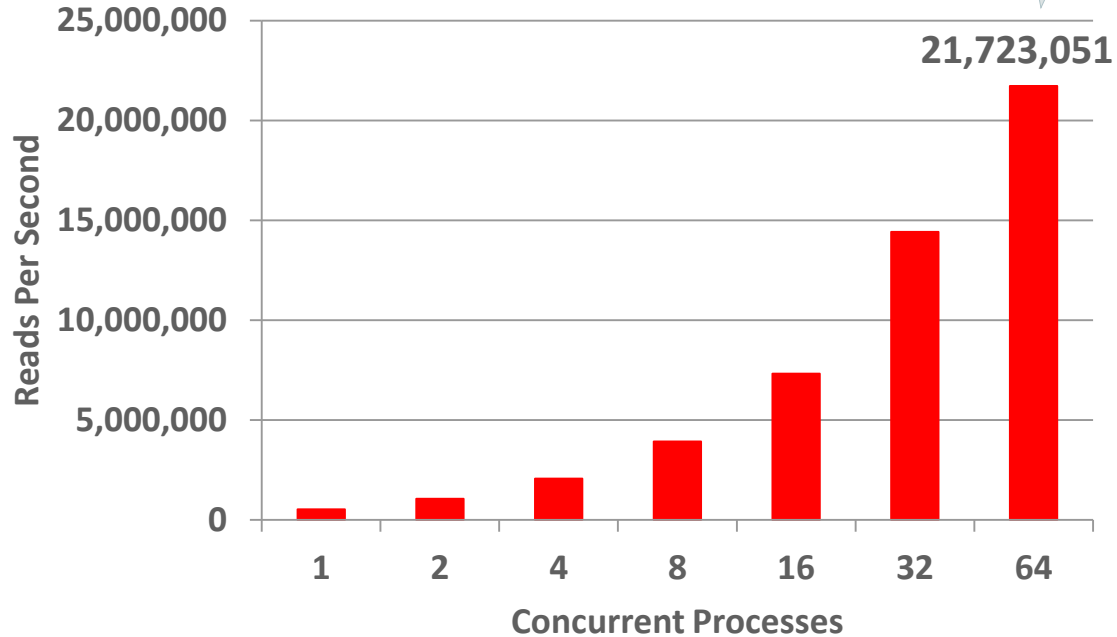
TPTBM 100% Mixed
Workload (80-10-5-5)
E5-2699 v4 @ 2.20GHz
2 socket, 22
cores/socket,
2 threads/core
TimesTen 11.2.2.8.0
(100M rows, 17GB)

TimesTen In-Memory Database

21.7 Million Transactions Per Second

**21.7 Million
Reads Per
Second**

Read-Only Workload Throughput

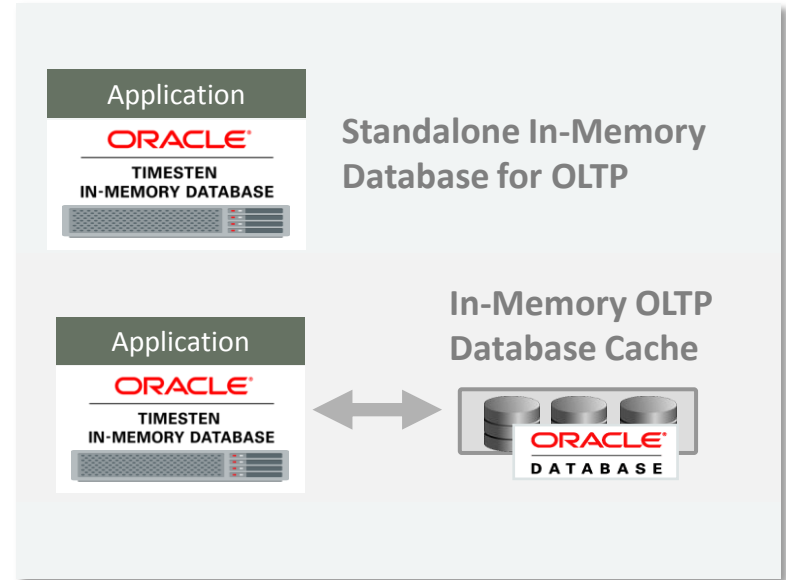


TPTBM 100% Read
E5-2699 v4 @ 2.20GHz
2 socket, 22
cores/socket,
2 threads/core
TimesTen 11.2.2.8.0
(100M rows, 17GB)

Oracle TimesTen In-Memory Database

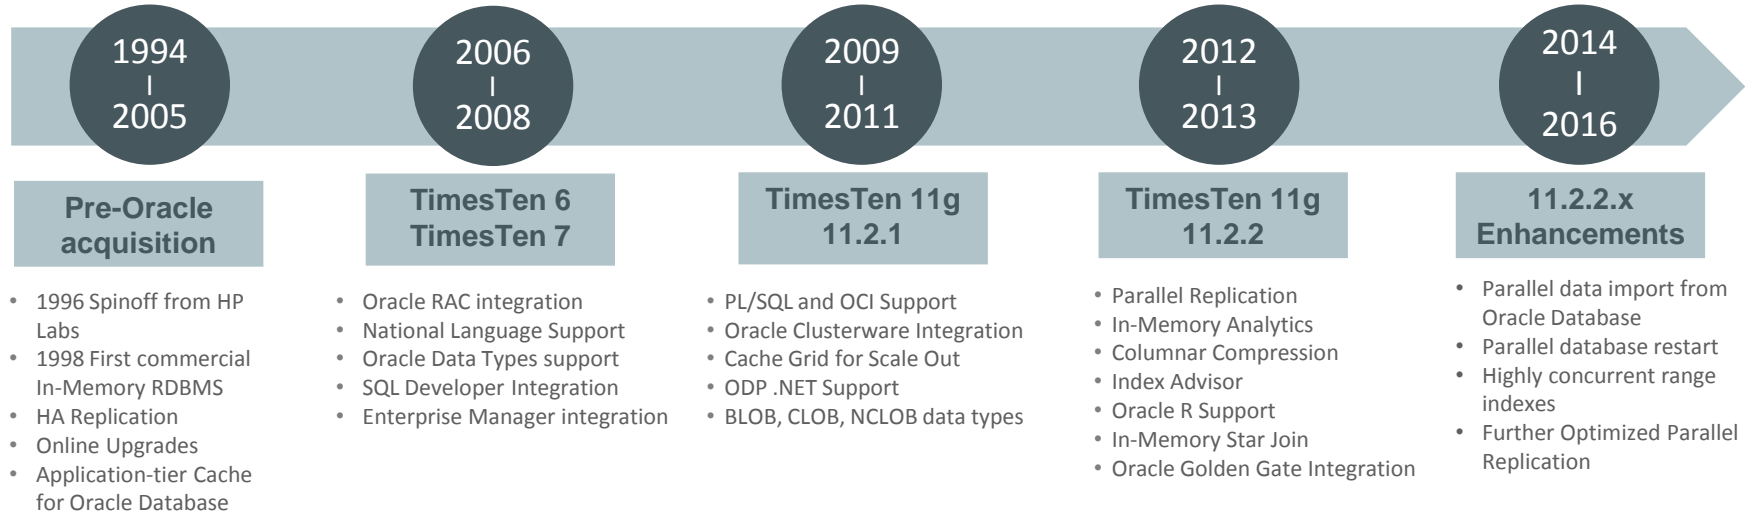
Two Deployment Options

1. Standalone In-Memory Database for OLTP applications
2. Application-tier Database Cache for the Oracle Database
 - Targeted for OLTP applications



Oracle TimesTen – Pure In-Memory Relational Database

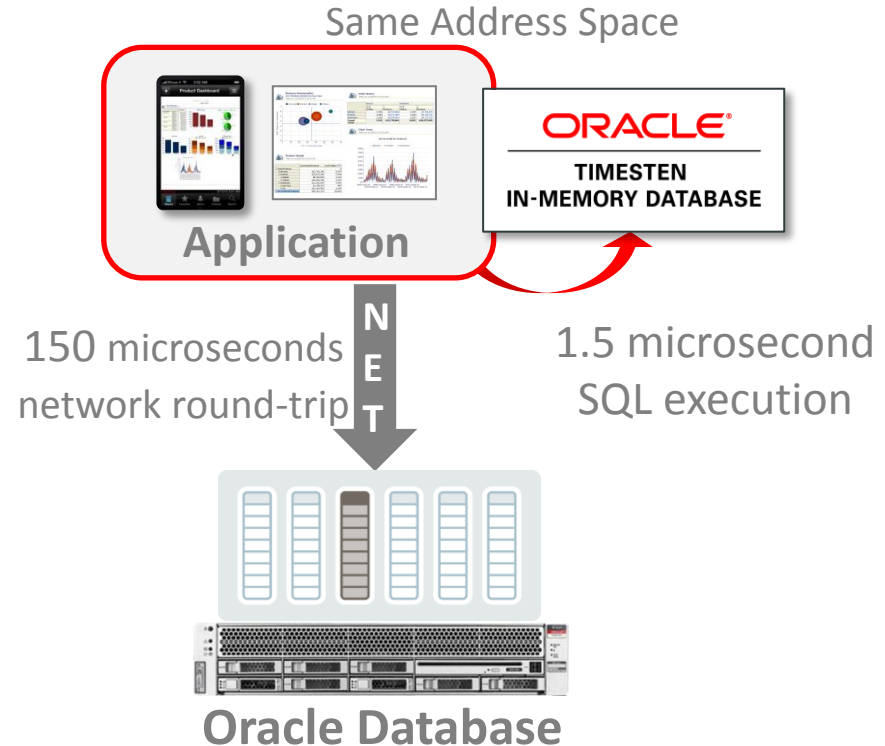
Mature In-Memory Technology



TimesTen: Primary Use Case

High Velocity Applications

- What are High Velocity applications?
 - Event driven, no opportunity to batch
 - E.g. Stock trade, phone call, credit-card authorization, need to be processed **immediately**
 - Usually lightweight transactions – few rows, very high transaction volume
 - TimesTen ideal for these use cases
- TimesTen is light-weight & ultra-fast
 - Runs in application: no network required
 - **100X** faster latency-critical applications
 - 1000s of customers



Most Widely Used In-Memory Database for OLTP Deployed by Thousands of Companies



Global Broker-Dealer

Mutual Fund Trading

- Program trading application serving institutional clients
- Business challenges
 - J2EE caching of full objects was too slow
 - Homegrown Java object cache too expensive for in-house staff
- TimesTen deployment (since 2003)
 - Order transaction processing
 - Pre-trade validation, order preparation and release, post-trade allocation
 - Reference data lookup
 - Event publishing to back-end database
- Why TimesTen
 - Standards based, commercial product
 - Order of magnitude performance improvement
 - Can sustain high volume orders

Real-Time Fraud Detection

USPS – Total Revenue Protection (TRP)



Challenges

- 4 billion mail scans per day peak (74,000/sec)
- 275 processing and distribution centers
- 33,000 postal facilities
- Find, track, and reject mail due to duplicate postage, short pay, or ineligible discounts
- 509 row inserts/sec (RIPS) → 275M txs per 15 hr processing window
- Sorting and capture time exceeded processing window

Solution

- Real-time data scans ingested into TimesTen
- 1.6TB TimesTen in-memory database
- Real-time TRP algorithms executed on TimesTen
- Results retained in TimesTen and propagated to Oracle Database for long term storage and analysis

TimesTen Values

- 190,222 RIPS (3 threads)
- 1,091,018 RIPS (18 threads)
- Processed 4 Billion txs in less than 6 hours
- Revenue protection performed in real-time upon first scan
- Sorting and capture easily fit within processing window

Mobile Phone Charging System

Ericsson Sweden



Challenges

- 5 Billion subscriptions in the world, 20% are charged via Ericsson
- Real-time Rating (price calculation, promotion and loyalty)
- Real-time accounting (spending control, multi-account and units, historical usage)
- Telecom grade, 99.999% availability, quick and automatic failover



Solution

- TimesTen In-Memory Database
- TimesTen Replication
- Shared nothing clusters
- Standard SQL interface
- Low maintenance
- Wide platform support
- Low system impact



TimesTen Values

- Predictable response time
- Very fast SQL performance
- High availability 99.999% up time (max down time 5 minutes per year)


AGENDA

1. TimesTen Introduction and Performance
2. TimesTen Velocity Scale
3. Q&A

What is TimesTen Velocity Scale In-Memory Database?

Scale-Out In-Memory Database for OLTP Applications

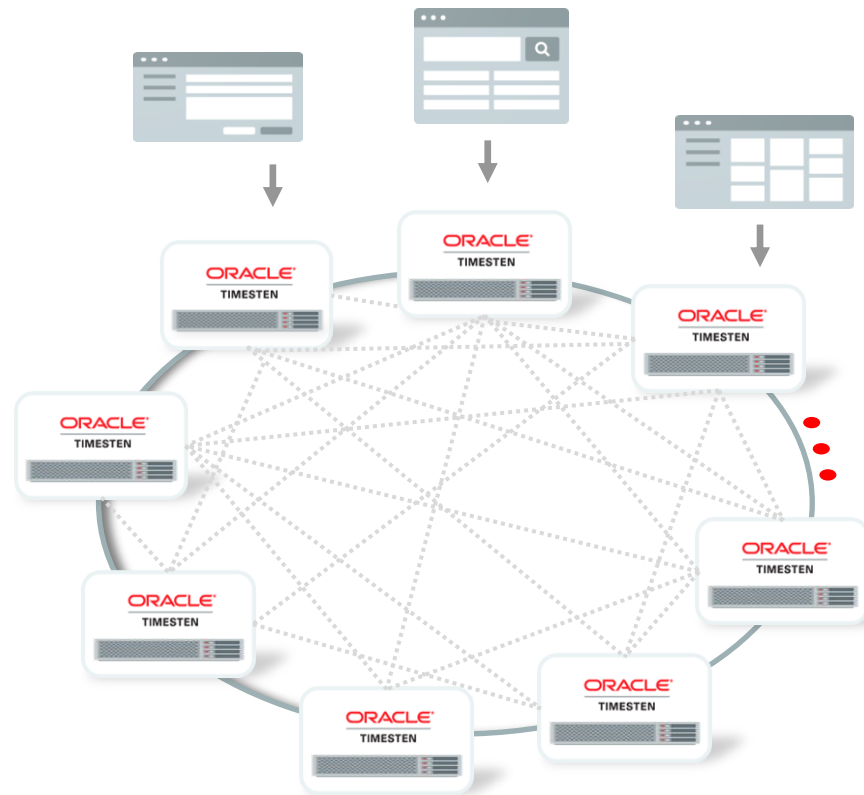
- An enterprise grade, distributed, highly available, shared nothing, scale-out, in-memory database
 - Based on TimesTen in-memory database technology
 - SQL relational
 - ACID compliant
 - Persistent and recoverable
 - Transactional
 - Scale-out architecture with single database image
 - Built-in high availability via multiple copies of the data
 - Elastic scalability
 - Easy to deploy, use and write apps for
- Designed for OLTP applications which require
 - Horizontal Scalability, High Throughput, High Availability
 - RDBMS semantics, consistency, functionality



Single Database Image

Connect to Any Host – Access ALL data

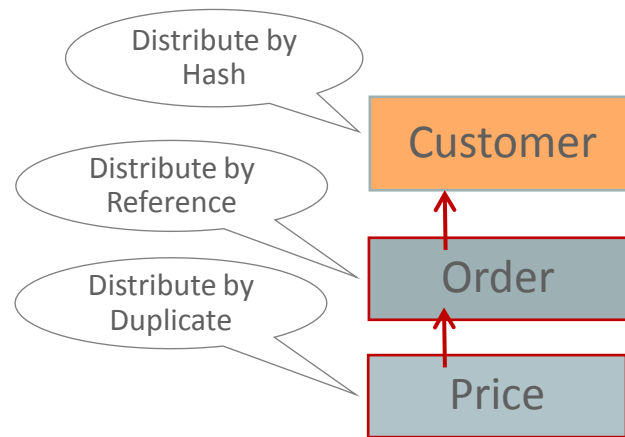
- Data distributed to all hosts
- Connect to any host and access all data
 - Execute queries and DML targeting data residing in any element(s)
 - Distributed queries, joins & transactions
- Concurrent transactions and parallel queries across all hosts
- All transactions may access / modify *any* data
 - Even across multiple hosts



Data Distribution Methods

Distribute by Hash and by Reference

- 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 for joins
 - Based on FK relationship
 - Supports multi-level hierarchy
 - Customer ← Order ← Line Items
- Distribute by Duplicate
 - Identical copies on all elements
 - Useful for reference tables
 - No remote access costs for reads
 - Join optimization



Host 1		Host 2		Host 3		Host 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		

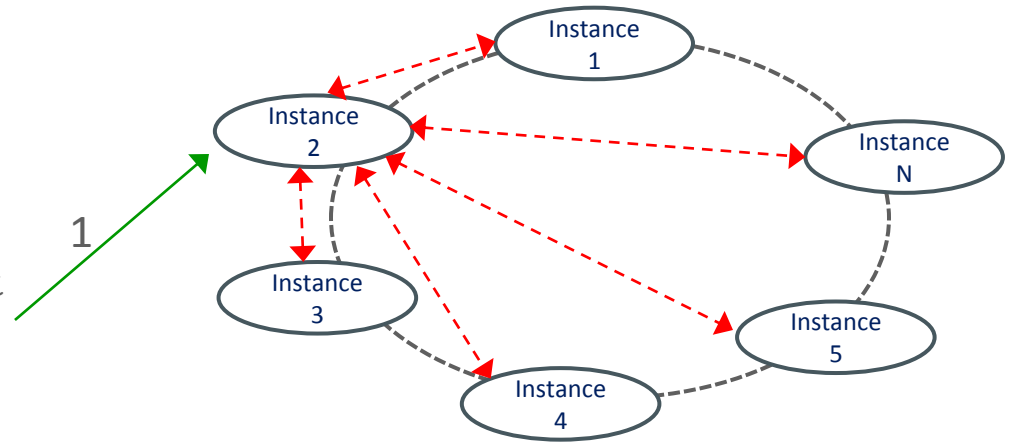
Parallel Query Processing

- Application transparent parallel query processing
 - Joins data across elements
 - Aggregate results

Example

```
select cust.name from cust
where cust.balance < 100;
```

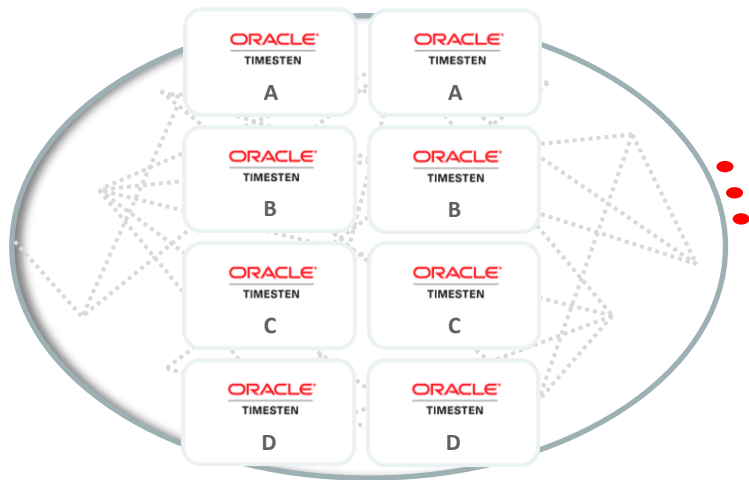
- One request from application
 - Result set aggregated by TimesTen Velocity Scale



— Application request
- - - Operations on behalf of app

High Availability and Maximum Throughput

K-Safety, All Active

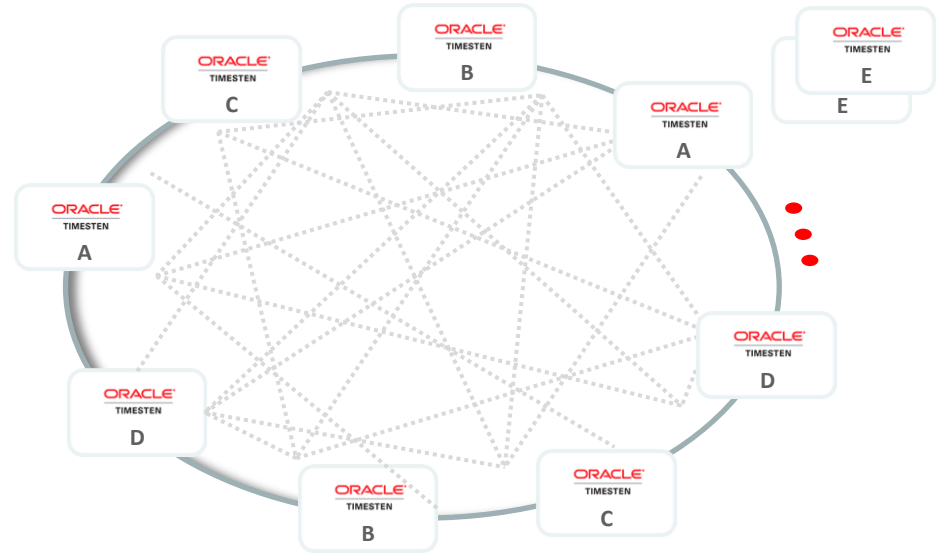


- Built-in HA via multiple copies of the data (K-safety)
 - Automatically kept in sync
- All replicas are active for reads and writes
- Transactions can be initiated from and executed on any replica
- Transactions may access / modify **any** data
 - Even across multiple hosts

Elastic Scalability

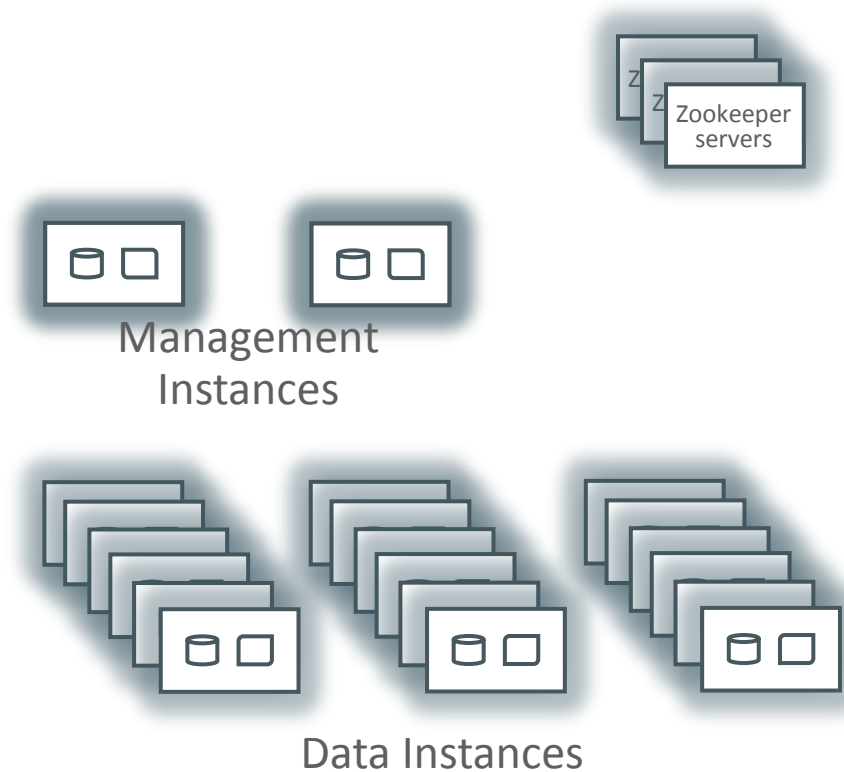
Expand the Database with Business Growth

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



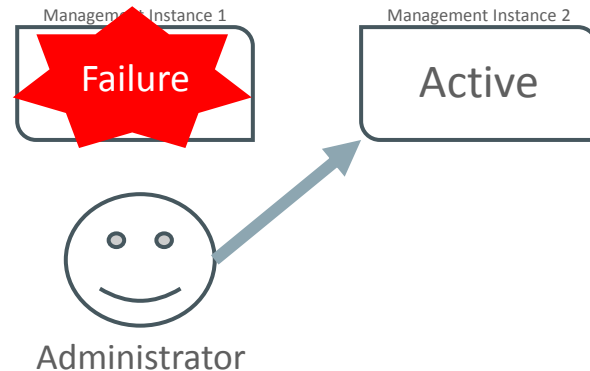
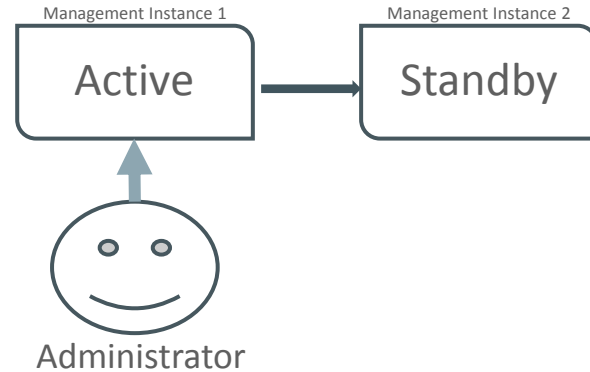
Grid Structure

- Each grid contains:
 - 1 or 2 management instances
 - A number of data instances
 - No fixed maximum
- Each grid uses:
 - A set of membership servers running Zookeeper (typically 1 or 3)
 - Can be shared by several grids



Grid management

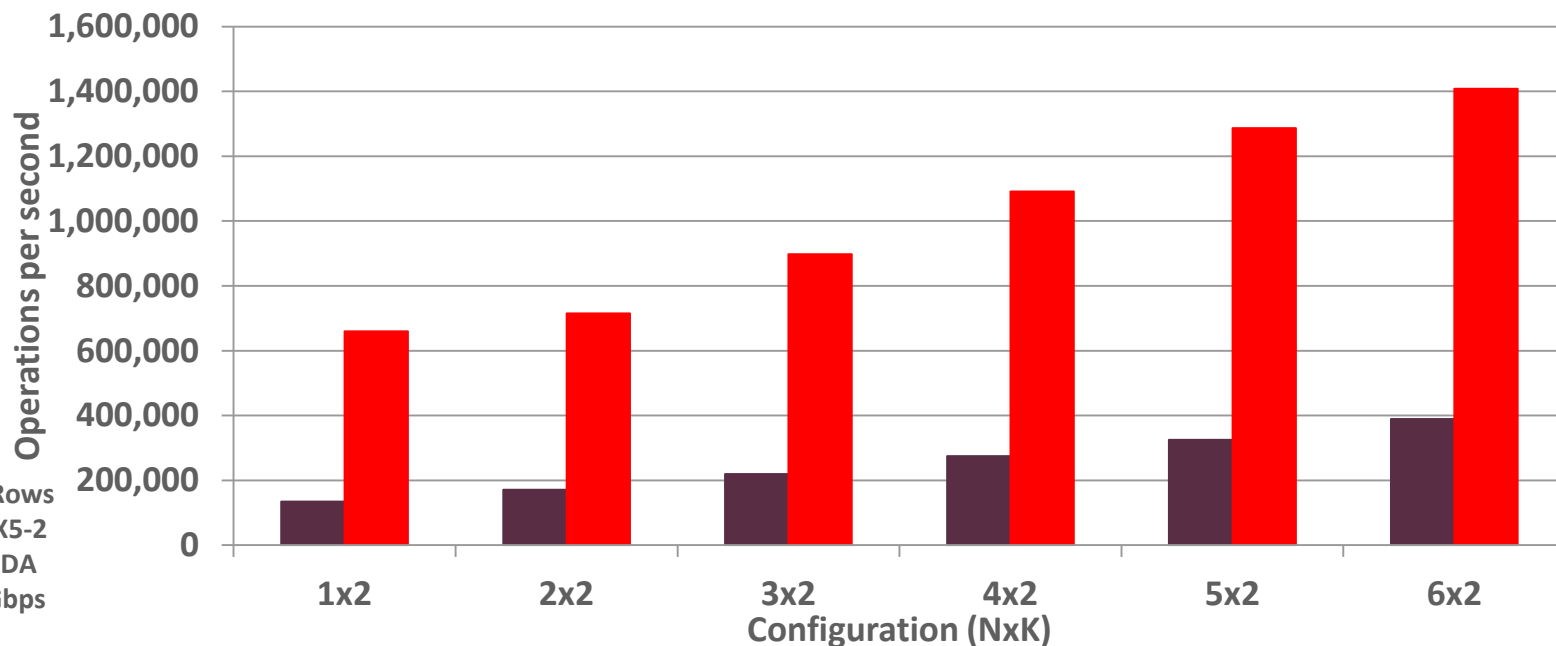
- Management instances work in an active / standby configuration
- **All** installation, configuration, management and administration of the grid is done on the active management instance
- You **never** have to log on to or copy files to each host manually
- The standby management instance can become the active in case of failure
- Management via:
 - Command line
 - SQL Developer



TimesTen Velocity Scale IMDB vs Cassandra

YCSB (Yahoo Cloud Service Benchmark) – 95% reads 5% writes

■ Cassandra ■ TimesTen Velocity Scale

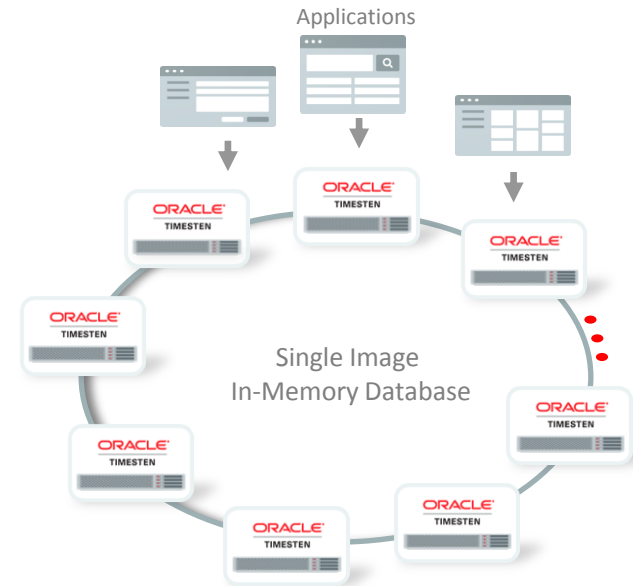


YCSB N*10M Rows
Servers: 12 x X5-2
Clients: 2 x ODA
Network: 10Gbps

What is TimesTen Velocity Scale In-Memory Database?

Scale-Out In-Memory Database for OLTP Applications

- An enterprise grade, distributed, highly available, shared nothing, scale-out, in-memory database
 - Based on TimesTen in-memory database technology
 - SQL relational
 - ACID compliant
 - Persistent and recoverable
 - Transactional
 - Scale-out architecture with single database image
 - Built-in high availability via multiple copies of the data
 - Elastic scalability
 - Easy to deploy, use and write apps for
- Designed for OLTP applications which require
 - Horizontal Scalability, High Throughput, High Availability
 - RDBMS semantics, consistency, functionality



AGENDA

1. TimesTen Introduction
2. TimesTen Velocity Scale
3. Q&A

ORACLE®