



Adding Speed and Scale to Oracle Database Deployments

Rob Meyer Outbound Product Management, GridGain

April 24, 2019

Glenn Wiebe Solution Architect, GridGain

2019 © GridGain Systems

Challenges with Speed and Scale *It's not just about the database*



GridGain Company Confidential

GridGain

Challenges with Speed and Scale It's not just about the database



GridGain Company Confidential

Simple Rules for Success with Speed and Scale

- 1. Ask the "5 why's" to find the speed and scale challenges
- 2. Optimize for end-to-end speed and scale
- 3. Look ahead over many projects
- 4. Choose the right long-term Oracle and 3rd party products



Oracle Database Options for Adding Speed and Scale



GridGain Company Confidential

GridGain

Oracle and 3rd Party Options

- Oracle Database In-Memory Option
- Oracle Database RAC
- Oracle Exadata
- Oracle TimesTen Classic
- Oracle TimesTen Application Cache
- Oracle TimesTen Scaleout
- Oracle Coherence



Oracle Database In-Memory Option Lowers latency in the database, doesn't add scale

- Lowers latency by adding a second, columnar in-memory version of an existing row-based table
 - Row table used for transactions, in-memory columnar table for analytics
 - Kept in synch with row disk-based table via ACID transactions
 - No additional work needed beyond configuration
- Benefit: only helps with latency, not scalability
- Cost: Additional 50% above enterprise edition license
 - Also requires additional RAM
- When to use: existing applications with slow query performance



Oracle Database RAC *"Somewhat" horizontal scalability and improved reliability*

- Oracle Real Application Clusters (RAC) mature 20+ year technology
 - "Shared something" architecture built on shared global cache that coordinates locks across servers in a RAC cluster
 - Generally used with fewer, bigger nodes (largest in 2017 was 48 nodes)
- Benefit: increases reliability, availability, write scalability
 - Adds reliability and availability compared to a single instance
- Cost: Additional 50% above enterprise edition license
 - Requires additional servers and more RAM per server to hold global cache
 - Can be lower-cost than a hardware-based approach for scaling up
- When to use: consider for most business-critical deployments





- Comes from Sun acquisition
 - X7-2, X7-8 systems available in 1/4-full rack configurations
 - Disk or flash-based storage
- Benefit: lowers latency, increases scalability
 - Vertical scalability
 - Can be used in combination with Oracle RAC
- Cost: \$1M or more for a single rack
- When to use: when it works for the long term and you don't have a choice
 - Works well when all you cannot change the Oracle-based application
 - Identify whether it will scale enough or whether you'll also need something else to help handle loads



Oracle GoldenGate The easiest way to offload reads and ETL

- Acquired by Oracle in 2009
 - Uses change data capture (CDC) off the write-ahead log (WAL)
 - Adds less than 5% additional load to a database
 - Used for real-time data extractions for ETL and operational data stores (ODS)
- Viable solution for offloading reads
 - GridGain integrated with GoldenGate
 - Good option when you cannot intercept all writes via in-memory computing
- Benefit: Offloads majority of load (reads) from Oracle Database with minimal cost
 - Eventually consistent only, though usually sub-second consistency
- Cost: lower than Oracle (see price list)
- When to use: a viable options for offloading reads
 - Eventual consistency is OK
 - Cannot easily add an in-memory data grid as a read- and write-through cache



Oracle TimesTen (Classic) Application-Tier Database Cache Lowers latency one application at a time

- Acquired by Oracle in 2005
 - TimesTen originally a scale-up in-memory database (IMDB), renamed TimesTen Classic
 - Added an option to use TimesTen Classic as a remote Oracle Database cache
 - Eventual consistency: writes to local cache then sent to Oracle. Oracle can also "refresh" TimesTen Classic cache with updates to Oracle Database.
 - Offers local disk-based backup and recovery for restarts
- Benefit: Lowers latency and offloads reads, but does not help with scalability
 - Supports most Oracle commands as-is (PL/SQL, trigger results propagated later)
 - Data kept locally, in RAM for much lower latency
 - Can only scale vertically on a single instance
- Cost: 50% of Oracle Database Enterprise Edition license
 - Requires additional RAM to hold all relevant data
- When to use:
 - Eventual consistency is OK
 - Cannot easily add an in-memory data grid as a read- and write-through cache



Oracle TimesTen Scaleout *The best rip-and-replace option for Oracle Database deployments*

- Released mid-2018 (1.0 product)
 - True distributed scale-out IMDB with affinity-based (e.g. foreign key) partitioning
 - Best support for Oracle Database SQL, PL/SQL
 - Initial release supports up to 64 nodes with 1 backup node each (128 total)
 - In-memory only. It does not use disk for storage. It does have backups
 - Standalone database. It does not act as a cache with Oracle database
- Benefit: Lowers latency and provides horizontal scalability
 - The simplest transition from Oracle Database-based applications
 - Improves latency and scalability
- Cost: Same cost as Oracle Database Enterprise Edition
 - Also buying new hardware
- When to use:
 - For custom applications with a lot of PL/SQL that need low latency and scalability
 - Not good for collocating data and external code (Java, .NET, C++, ...)





GridGain Company Confidential

GridGain

How an In-Memory Data Grid Works

How an IMDG Works

- Slides in-between Oracle
 Database and the app / analytics
- Acts as the new database

Benefits

- No rip-and-replace of Oracle
- In-memory speed
- Horizontal scalability
- Collocated computing
- Merges data across sources
- Unlocks data for new applications



In-Memory Data Grid (IMDG)





GridGain

Oracle Coherence Not a good long term option

- Acquired by Oracle in 2007 (Tangasol)
 - Had most of the capabilities of an IMDG at the time
 - Has not evolved since
 - Independent release: 3.7 released in 2011, currently 3.7.1
 - Oracle Database bundle: 12.2.1.3
- When to use: keep for existing deployments
 - Is not a competitive option compared to other IMDGs

	FEATURE	GRIDGAIN CE 2.7 (APACHE IGNITE 2.7)	GRIDGAIN EE 8.5	GRIDGAIN UE 8.5	ORACLE COHERENCE 12.2.1 (3.7.1)
	Native ANSI-99 SQL Support	•	•	•	•
an	Distributed ACID Transaction Support	•	•	•	(No explicit locking, performance issues)
	Slide in-between SQL- based Applications and RDBMSs with with support for SQL	•	•	•	(Requires code for new apps, adding code and data model for existing apps)
	Cross-Language Support for Collected Processing (MPP)	(Supports Multiple Languages, MPP for Java, .NET, C++)	(Supports Multiple Languages, MPP for Java, .NET, C++)	(Supports Multiple Languages, MPP for Java, .NET, C++)	•
.3	Integration with RDBMSs, NoSQL Databases and Hadoop	(Out-of-the box support for RDBMSs, NoSQL Databases, HDFS, Spark)	(Out-of-the box support for RDBMSs, NoSQL Databases, HDFS, Spark)	(Out-of-the box support for RDBMSs, NoSQL Databases, HDFS, Spark)	(Requires coding. No Spark or HDFS support)
	Comprehensive In-Memory Computing Solution	(IMDG, streaming, machine and deep learning)	(IMDG, streaming, machine and deep learning)	(EE + Multi-datacenter data and disaster recovery management)	(IMDG only. No IMDB, streaming, machine or deep learning)
	Apache Spark Support for DataFrames, RDDs, HDFS	•	•	•	•
	Built-in Machine Learning	•	•	•	•
GridGain Co	Built on a Leading Open Source Project	(Built on Apache Ignite, a top 5 Apache Software Foundation	(Built on Apache Ignite, a top 5 Apache Software Foundation	(Built on Apache Ignite, a top 5 Apache Software Foundation	(No open source project or external development community)



Real-time Business Add speed and scale to existing applications

Store new types of (big) data

A Common Data Layer for

- Ingest and process streaming data with Apache Spark and other streaming analytics technologies to support real-time analytics
- Implement real-time decision automation including continuous machine and deep learning



GridGain Company Confidential

In-Memory Computing Comparisons https://www.gridgain.com/resources/product-comparisons

Vendors Compared

- Oracle Coherence
- Pivotal Gemfire (Apache Geode)
- GigaSpaces
- GridGain (Apache Ignite)
- Hazelcast
- Redis (Cache)
- Terracotta



GridGain In-Memory Computing Platform

- Built on Apache Ignite
 - Comprehensive platform that supports all projects
 - No rip and replace
 - In-memory speed, petabyte scale
 - Enables HTAP, streaming analytics and continuous learning
- What GridGain adds
 - Production-ready releases
 - Enterprise-grade integration, security, deployment and management
 - Global support and services
 - Proven for mission critical apps

GridGain In-Memory Computing Platform





Accelerate Existing Applications with No Rip and Replace



- Slides in-between apps and RDBMSs with no rip and replace
 - ANSI-99 SQL compliant
 - Support for ACID transactions
- Accelerates existing app performance
- Offload new data and computing requirements (real-time auditing and compliance, analytics, computations)





Innovate with Existing and New Data

GridGain as an In-Memory Database (IMDB)

Memory-centric storage

- From 100% in-memory to 100% disk
- Leverages any combination of RAM, Flash, SSD, Intel 3D Xpoint and disk
- Low cost, disk-based reliable persistence
- Immediate restart during recovery
- Highest read+write performance
 - In-memory with unlimited linear, scale-out on commodity servers
 - SQL and NoSQL (multi-model)
 - Always-on availability
- Single data access layer for ALL data
- Extensible compute grid



Innovate with Streaming Analytics

GridGain for Stream Ingestion, Processing and Analytics

- Native support for stream ingestion
 - Built-in support for high speed ingestion from Apache Camel, Flink, Flume, Spark, Storm, JMS, Kafka and MQTT
 - Combines streams with data-at-rest
 - Collocated data processing across all data, including optimized SQL querying
 - Publish/subscribe (continuous queries)
- Broadest in-memory support for Apache Spark
 - Native in-memory RDD, DataFrame support
 - Shares state in memory across Spark jobs
 - Native access to ANY data across GridGain cluster
 - Optimizes SparkSQL using distributed SQL and indexing



Innovate with Continuous Learning

Continuous Learning Framework for Machine and Deep Learning

- Real-time performance on petabytes of data
 - No ETL (runs learning in place)
 - In-memory performance
 - Horizontal, linear scalability
- Machine learning
 - Linear, multi-linear regression
 - K-means clustering
 - Decision trees
 - K-NN classification and regression
- Deep Learning
 - TensorFlow integration



Wellington - Next Generation, Real-time IBOR

A top 20 worldwide asset management firm with over \$1 trillion under management

Problem

- Current systems no longer scaled to handle the volumes
- Didn't comply with new regulations following financial crisis
- Needed to introduce new asset classes faster

GridGain Solution

- Investment Book of Record (IBOR), a single real-time version of the truth for positions, exposure, valuations and performance for all customers, teams and trades.
- 10x performance gains, linear horizontal scalability
- Support for SQL and ACID transactions, and for existing systems and skillsets
- Enabled transactions and analytics on a single platform
- Collocated computing scales complex calculations, analytics

W E L L I N G T O N MANAG E M E N T®





Demo

How to Configure GridGain as an IMDG for Oracle



Questions?





In-Memory Computing Summit Europe Starts June 3rd in London



Enter Weekly Raffle for a Chance to Win a £20 Ticket Voucher

https://www.imcsummit.org/2019/eu/imcs-eu2019-raffle



GridGain Resources https://www.gridgain.com/

Ignite Resources https://ignite.apache.org/

- Webinars
 - GridGain: https://www.gridgain.com/resources/webinars
 - In-Memory Computing Summit: <u>https://www.imcsummit.org/</u>
- White Papers: https://www.gridgain.com/resources/papers
- Videos: <u>https://www.gridgain.com/resources/videos</u>
- Blogs: <u>https://www.gridgain.com/resources/blog</u>
- Downloads
 - Apache Ignite: https://ignite.apache.org/download.cgi
 - Free GridGain Community Edition, or free 30-Day Enterprise or Ultimate Edition Trial

https://www.gridgain.com/resources/download







Adding Speed and Scale to Oracle Database Deployments

Rob Meyer Outbound Product Management, GridGain

April 24, 2019

Glenn Wiebe Solution Architect, GridGain

2019 © GridGain Systems