



# Adding Speed and Scale to PostgreSQL Database Deployments

Rob Meyer Outbound Product Management, GridGain

May 22, 2019

2019 © GridGain Systems

# **History of PostgreSQL**

- Leading Open Source RDBMS
  - Ever since MySQL acquisition by Sun/Oracle ...
- **Over 30 Different Variants** 
  - Host of OLTP and OLAP
  - Cloud offerings as well
- More Recent Focus on Scale
  - 9.0 (2015): Replication
  - 10.0 (2017): Logical replication
  - 11.0 (2018): Improvements to partitioning and parallelism

2000

2010





https://wiki.postgresgl.org/wiki/PostgreSQL derived databases

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





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



- 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
- Choose the right long-term PostgreSQL-related and 3<sup>rd</sup> party products



# **PostgreSQL Database Options for Adding Speed and Scale**



2019 © GridGain Systems

GridGain Company Confidential

GridGain

# **PostgreSQL** and 3<sup>rd</sup> Party Options

- Options to Lower Latency
  - Pgmemcache
  - Column store extension
- Options to Improve Scalability (OLTP)
  - PostrgreSQL replication
  - Postrgres-XL
  - 2ndQuadrant
  - EnterpriseDB
  - Citus Data
  - Amazon RDB for Postgres
- Options to Improve Scalability (OLAP) ...



### In-Memory Column Store and Pgmemcache Lowers read latency but not at scale

- "Built-in" Options
  - Use a RAM disk for storage
  - UNLOGGED tables in RAM
- In-Memory Column Store
  - Drop-in extension stores data as columns in shared buffers (RAM)
  - Slows down transactions, suffers from instabilities, potential data loss
- Pgmemcache
  - User-defined functions to use memcache in the context of PostgreSQL
  - Same limitations as memcache: cache aside cache that lowers latency but must be managed by application
  - Doesn't help with scalability, so doesn't solve limited network bandwidth for large data sets



# How to Scale PostgreSQL Horizontally

- Replication for Read Scalability
  - PostgreSQL Replication: WAL-based replication using streaming (hot and warm)
  - Manual failover
  - No automatic recreation of secondary
- Sharding for Write Scalability
  - A manual process
  - Hard to manage sharding to balance different workloads
- Cloud Offerings are One Option
  - Amazon RDS and Aurora
  - Azure and Google offerings
- OLAP Offerings Exist as Well, But OLAP and OLTP Are Separate
  - HP Vertica
  - Teradata (Aster Data)
  - IBM Netezza
  - Pivotal Greenplum (now open source)



### Pgpool-ii – Load Balancing Middleware https://www.pgpool.net

- Open Source (BSD License)
- Middleware Load Balancer (transparent proxy)
  - Connection pooling
  - Replication
  - Load balancing
  - Queues connections (versus rejecting them)
  - Query caching (lowers latency by caching SELECT statements)
- Good for simplifying load balancing and connection management to improve read scalability
- Not recommended by EnterpriseDB for use as the replication mechanism



# EnterpriseDB – Heterogeneous Replication <a href="https://www.enterprisedb.com">https://www.enterprisedb.com</a>

#### • Founded in 2004

- Have a larger EDB Postgres Platform
- Employ several of the the leading PostgreSQL contributors
- EDB Postgres Replication Server (EPRS)
  - Single-master or multi-master (bi-directional) replication
  - Also supports Oracle and SQL Server (using non-XA triggers)
- The only out-of-the-box heterogeneous replication solution
- Not Well Suited For:
  - Immediate consistency: relies on async, not transactional replication
  - Performance: trigger-based replication for Oracle, SQL Server adds significant load.
    Also, Infinite Cache (read/write-through cache) deprecated as of Release 8.2
  - Scalability: need to manage partitioning, rebalancing manually



### 2ndQuadrant – Distributed Replication (Postgres-BDR) https://www.2ndquadrant.com/

- Founded in 2001 by a Leading PostgreSQL Contributor
- Postgres-BDR (Bi-Directional Replication)
  - Open source
  - PostgreSQL only
  - Bi-directional replication with immediate or eventual consistency
  - Supports replication and sharding
  - Supports pushdown to shards with aggregation to minimize network traffic
- Not Well Suited For:
  - Heterogeneous database workloads
  - Implementing a distributed database ...



### 2ndQuadrant – Distributed PostgreSQL (Postgres-XL) https://www.2ndquadrant.com/

- Another Offering from 2<sup>nd</sup> Quadrant
  - Postgres-BDR: read/write replication
  - Postgres-XL: distributed PostgreSQL
- Postgres-XL
  - Replication to scale reads
  - Dynamic data redistribution and balancing
  - Distributed transaction processing (with an XA coordinator)
- Not Well Suited For:
  - Standard PostgreSQL (a fork that isn't 100% compatible)
  - High availability: takes time for a warm node to come up
  - Only tested for asynchronous transactions, uses one core per node



### Citus Data – Distributed PostgreSQL (Citus) https://www.citusdata.com/

- Acquired by Microsoft January 2019
  - Open source PostgreSQL extension
  - Also Enterprise version on-premise, Azure/AWS SaaS
- Citus is the Best Distributed PostgreSQL
  - Coordinator node manages distributed SQL across worker nodes
  - Supports zero-downtime for rebalancing shards as nodes are added
  - As of release 7.1 (2017) supports synchronous distributed transactions
- Not Well Suited For:
  - Lowering latency: same limitations as PostgreSQL on each worker node
  - Heterogeneous database environments: PostgreSQL-only
  - Java, .NET, C++ processing: only distributed SQL.





GridGain Company Confidential

GridGain

# How an In-Memory Data Grid Works

### How an IMDG Works

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

### Benefits

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



### In-Memory Data Grid (IMDG)



### 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





#### Add speed and scale to existing applications Store new types of (big) data

Ingest and process streaming data with Apache Spark and other streaming analytics technologies to support real-time analytics

**Real-time Business** 

Implement real-time decision automation including continuous machine and deep learning



### 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

In-Memory Data Grid	- - - 	n-Memory Database	Streaming Analytics		( Learn	Continuous hing Framework	
Machine and Deep Learning							
Messaging		Stream Processing			Events		
Key-Value		ANSI-99 SQL			ACID Transactions		
Compute and Service Grid							
In-Memory Data Store							
Persistent Store							
Jainframe F	RDBIMS	Data	Layer	NoSC		Hadoop	

# **Case Study: FSB Sports Betting Platform**

https://www.imcsummit.org/2018/eu/session/memory-computing-and-sports-betting





A leading sports betting platform as a service that supports thousands of bets per second processed across 750+ casino and live dealer games across 36 branded sites, 11 countries, 4 continents

#### Challenges

- Real-time ingestion of betting data
- Real-time calculation of betting odds and liability
- Lack of SQL querying with Memcached

#### **GridGain Enables**

- Real-time transactions and analytics
- Dynamic scalability across hybrid cloud
- Multi-datacenter availability





# **Questions?**





### Driving In-Memory Computing Thought Leadership GridGain produces the world's only In-Memory Computing Conference

### https://www.imcsummit.org/

In-Memory Computing Summit **Europe** 

• Next event in London, June 2019

In-Memory Computing S U M M I T

In-Memory Computing Summit North America



• Next event in Silicon Valley, Nov. 2019



### 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: <a href="https://www.gridgain.com/resources/papers">https://www.gridgain.com/resources/papers</a>
- 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 PostgreSQL Database Deployments

Rob Meyer Outbound Product Management, GridGain

May 22, 2019

2019 © GridGain Systems