ORACLE®

Using MySQL for Big Data Advantage

Integrate for Insight

Sastry Vedantam
sastry.vedantam@oracle.com

Agenda

- The rise of Big Data & Hadoop
- MySQL in the Big Data Lifecycle
- MySQL Solutions for Big Data
- Q&A

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 decision. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

DRIVING MySQL INNOVATION

MySQL Enterprise Monitor 2.2
MySQL Cluster 7.1
MySQL Cluster Manager 1.0
MySQL Workbench 5.2
MySQL Database 5.5
MySQL Enterprise Backup 3.5
MySQL Enterprise Monitor 2.3
MySQL Cluster Manager 1.1

All GA!

2010

MySQL Enterprise Backup 3.7
Oracle VM Template for MySQL Enterprise Edition
MySQL Enterprise Oracle Certifications
MySQL Windows Installer
MySQL Enterprise Security
MySQL Enterprise Scalability

All GA!

MySQL Database 5.6 DMR*
MySQL Cluster 7.2 DMR

MySQL Labs!
("early and often")

2011

MySQL Cluster 7.2
MySQL Cluster Manager 1.4
MySQL Utilities 1.0.6
MySQL Migration Wizard
MySQL Enterprise Backup 3.9
MySQL Enterprise Audit
MySQL Database 5.6
MySQL Cluster 7.3

All GA!

MySQL Database 5.7.2 DMR

A BETTER MySQL

2012-13

ORACLE

*Development Milestone Release

Pluggable Storage Engines Architecture

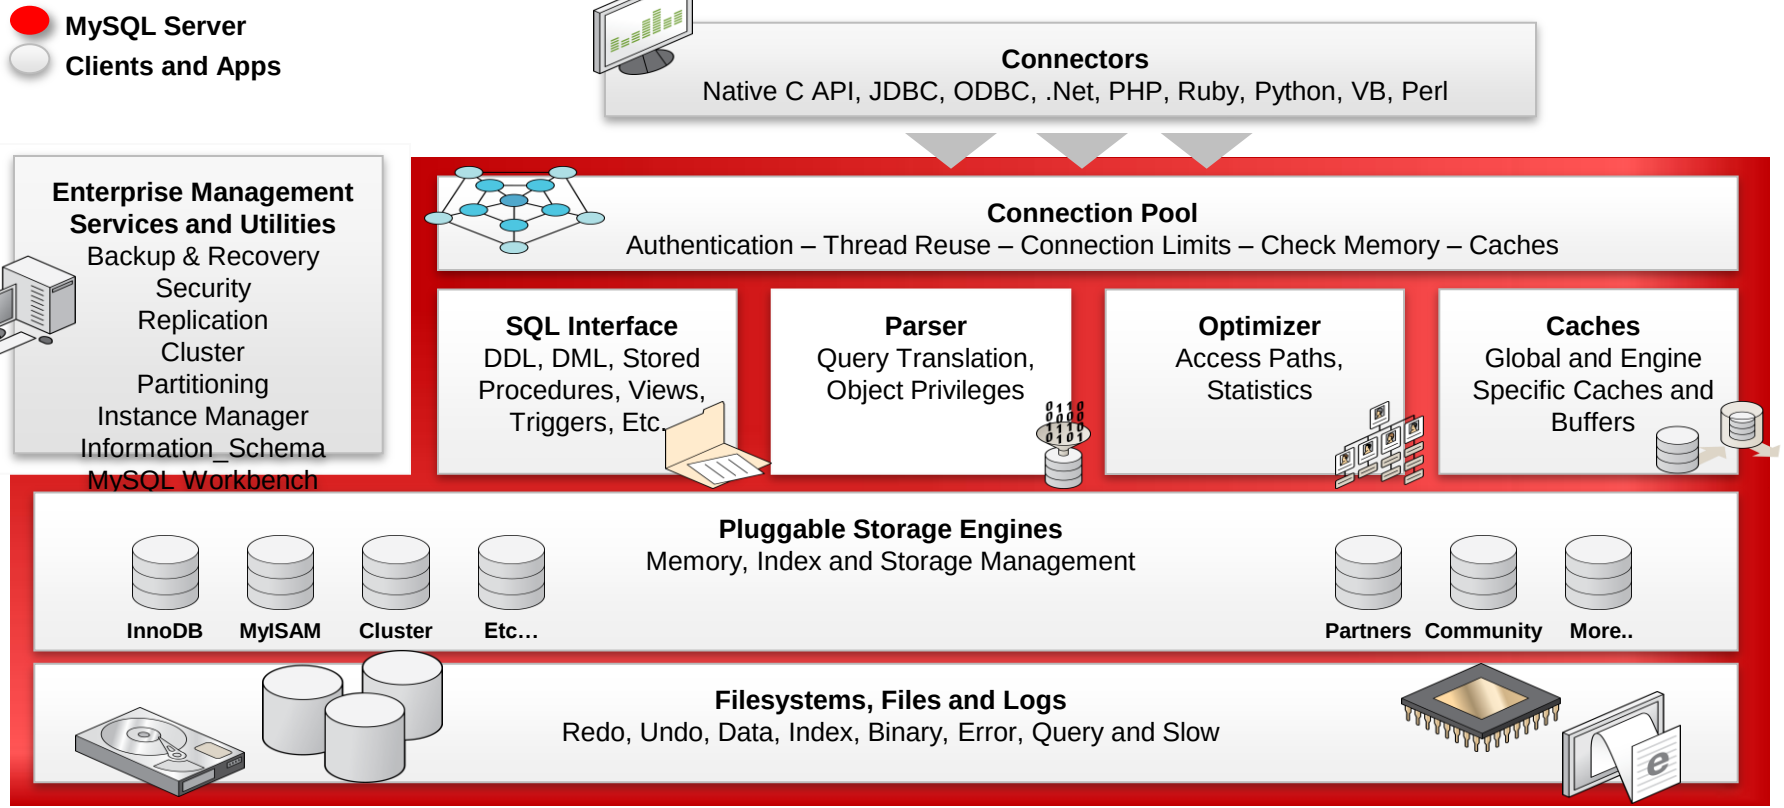

Industry Leaders Rely on MySQL

MySQL 5.6: In Summary

IMPROVED PERFORMANCE AND SCALABILITY

- Scales to 48 CPU Threads
- Up to 230% performance gain over MySQL 5.5

IMPROVED INNODB

- Better transactional throughput and availability

IMPROVED OPTIMIZER

- Better query exec times and diagnostics for query tuning and debugging

IMPROVED REPLICATION

- Higher performance, availability and data integrity

IMPROVED PERFORMANCE SCHEMA

- Better Instrumentation, User/Application level statistics and monitoring

New! NoSQL ACCESS TO INNODB

- Fast, Key Value access with full ACID compliance, better developer agility

MySQL 5.6: Best Replication Features Ever

PERFORMANCE

- Multi-Threaded Slaves
- Binary Log Group Commit
- Optimized Row-Based Replication

FAILOVER & RECOVERY

- Global Transaction Identifiers
- Replication Failover & Admin Utilities
- Crash Safe Slaves

DATA INTEGRITY

- Replication Event Checksums

DEV/OPS AGILITY

- Time Delayed Replication
- Remote Binlog Backup
- Informational Log Events

Leading Use-Case, On-Line Retail

Users

Browsing

Recommendations

MySQL

Profile,
Purchase
History

Recommendations

Social media updates
Preferences
Brands “Liked”

Web Logs:
Pages Viewed
Comments Posted

Telephony Stream

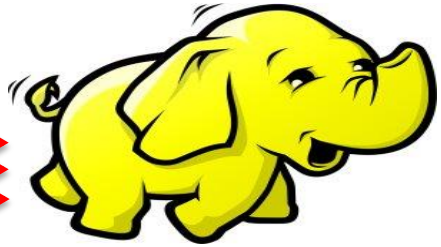

MySQL in the Big Data Lifecycle

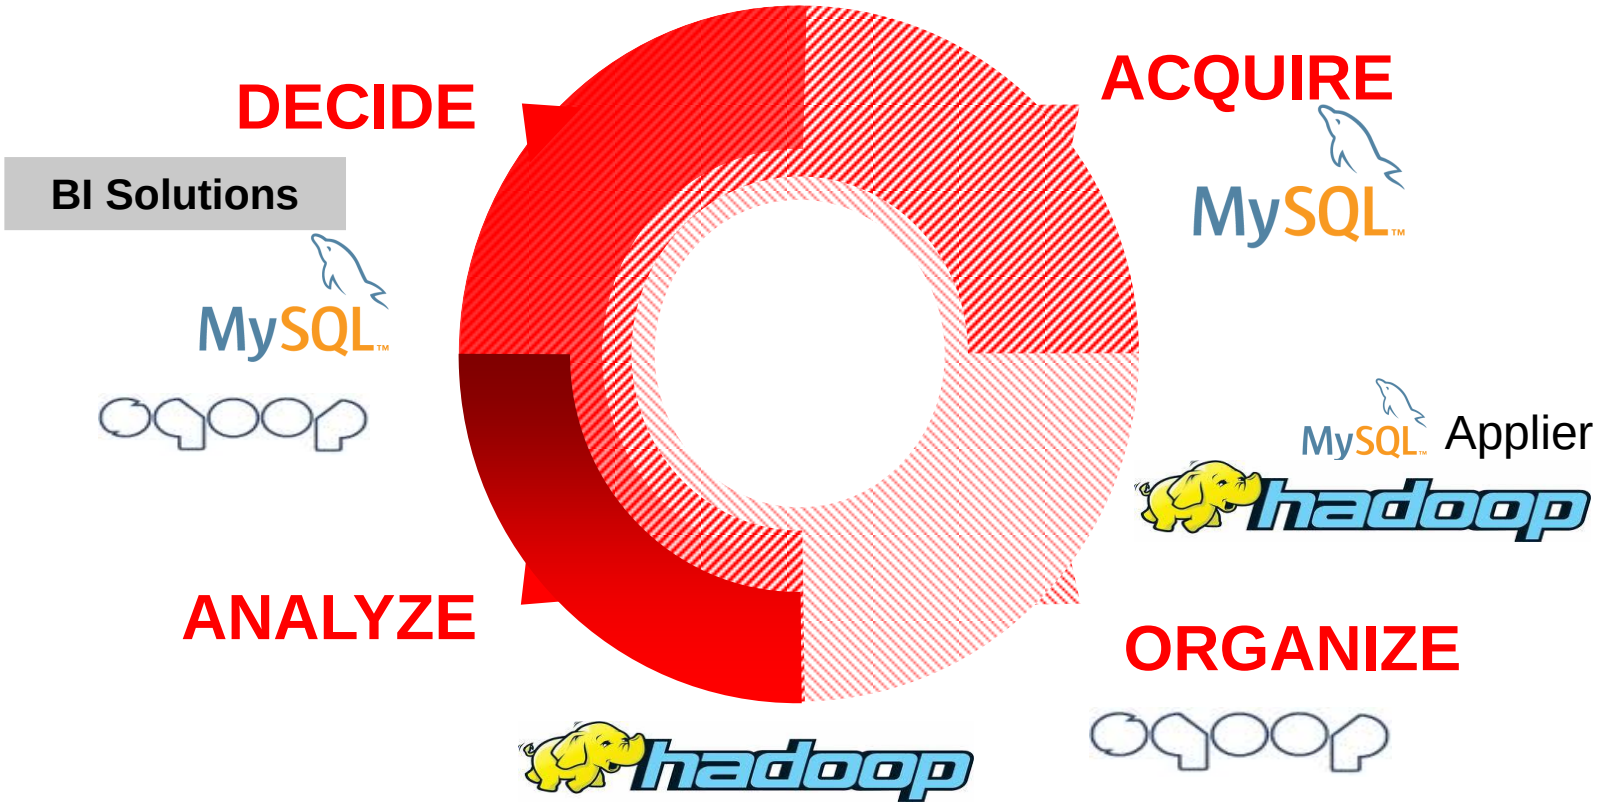

MySQL + Hadoop:
Unlocking the Power of Big Data

*50% of our users integrate with MySQL**

Download the MySQL Guide to Big Data:

http://www.mysql.com/why-mysql/white-papers/mysql-and-hadoop-guide-to-big-data-integration/

*Leading Hadoop Vendor

ORACLE

MySQL in the Big Data Lifecycle

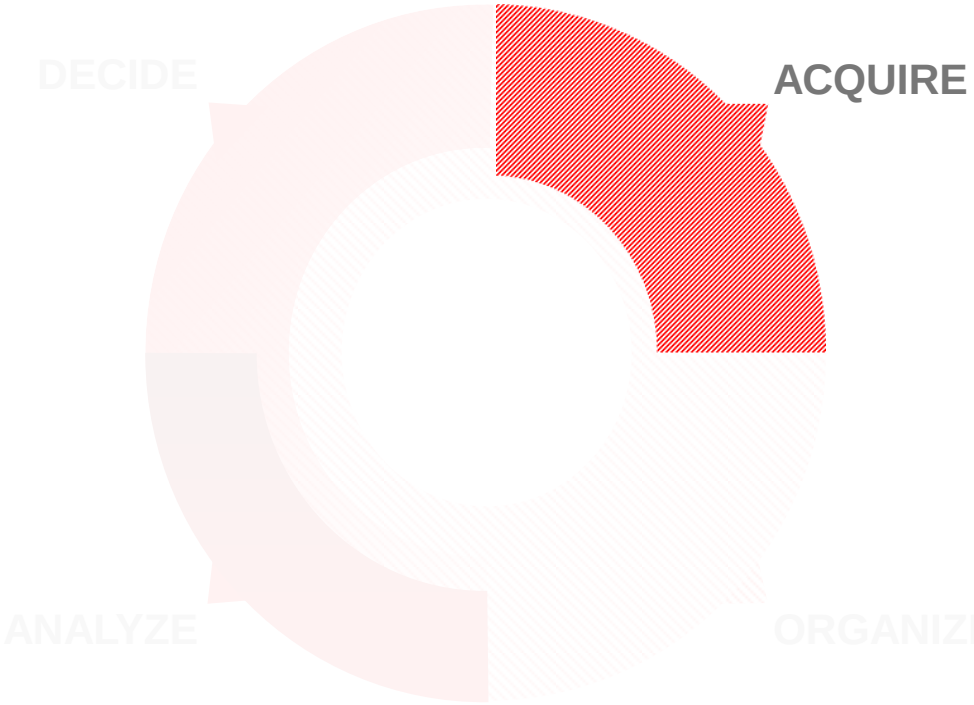

NoSQL Interfaces for

MySQL Database

MySQL Cluster

MySQL NoSQL Interface

Design Goals: Fast, Flexible and Safe

Blazing Fast Key / Value Queries

Fully transactional / ACID

NoSQL + SQL across same Data Set

Combined with Schema Flexibility: Online DDL

How Memcached is used with MySQL separately

- Memcached is in-memory key-value store for small data
- It is one of the most widely used In-Memory cache implementations for social network websites
- Memcached has a simple and open protocol as opposed to a rich client bound to a specific language, and implementation makes it portable across a wide variety of languages and environments

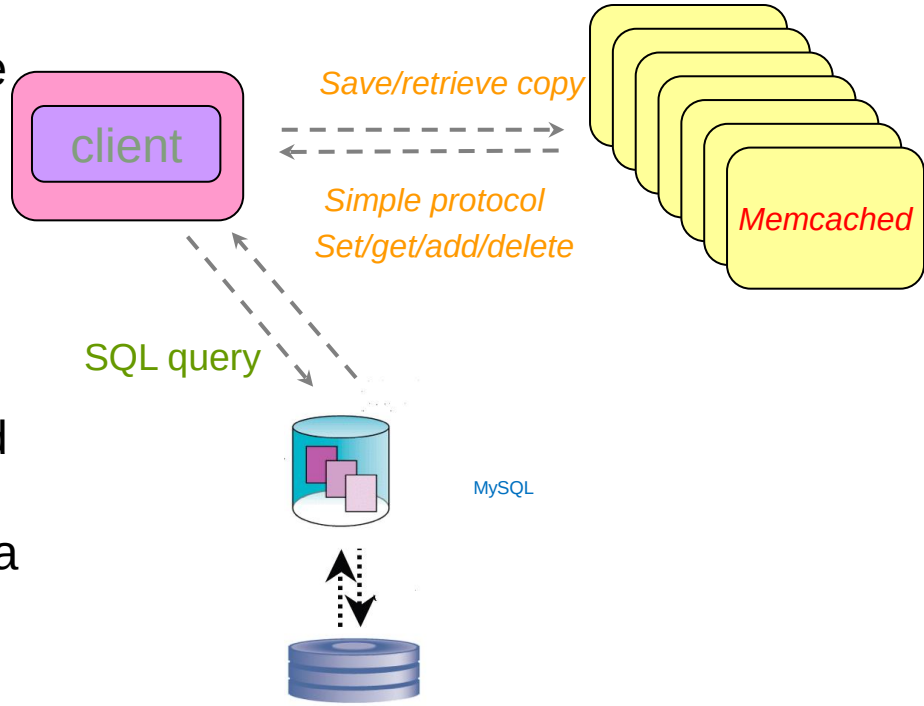

InnoDB as a Key Value store

- Combine the best of the NoSQL world and SQL world
- Memcached listens on specific ports as the front end, directs requests directly to InnoDB
- Simple commands, much smaller network transmit packages
- Persistent storage from InnoDB
- Index on the key column
- Full ACID compliance
- Bypass Optimizer and QP layer of MySQL and directly access the storage engine
- Dual access of data (SQL and Memcached)

MySQL 5.6: NoSQL Interface to InnoDB

Memcached API

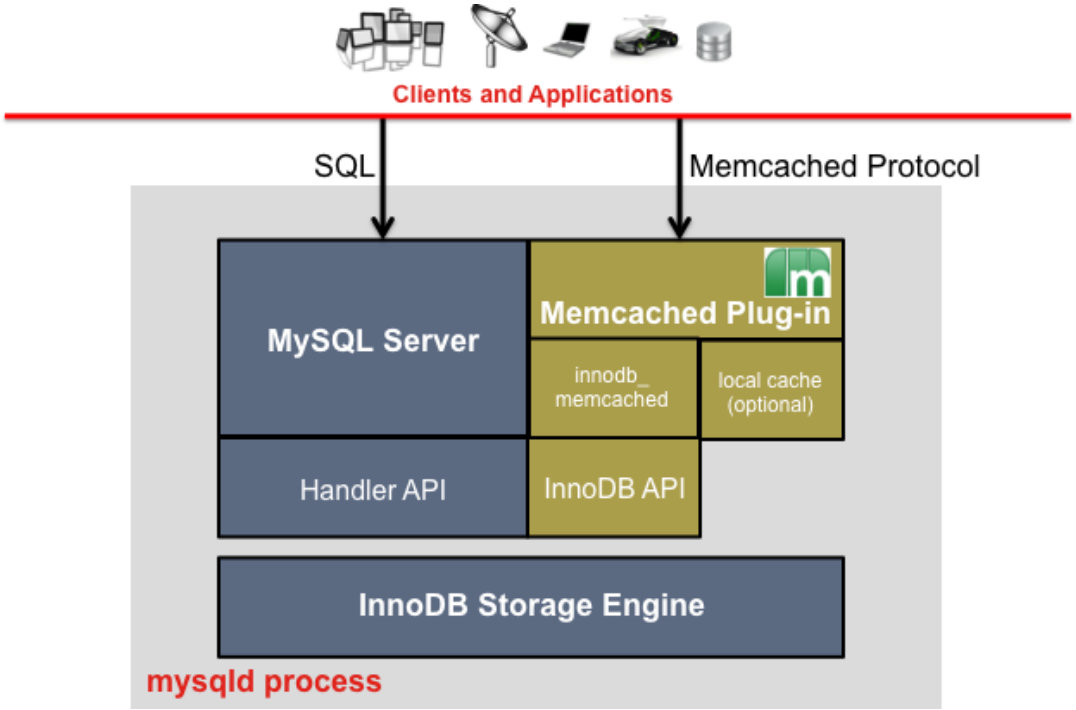

- **Key-value access to InnoDB**
- **Bypasses SQL parsing**
- **Implemented via:**
  - Memcached plug-in to mysqld
  - Memcached mapped to native InnoDB API
  - Use existing Memcached clients
  - Shared process for ultra-low latency

Performance

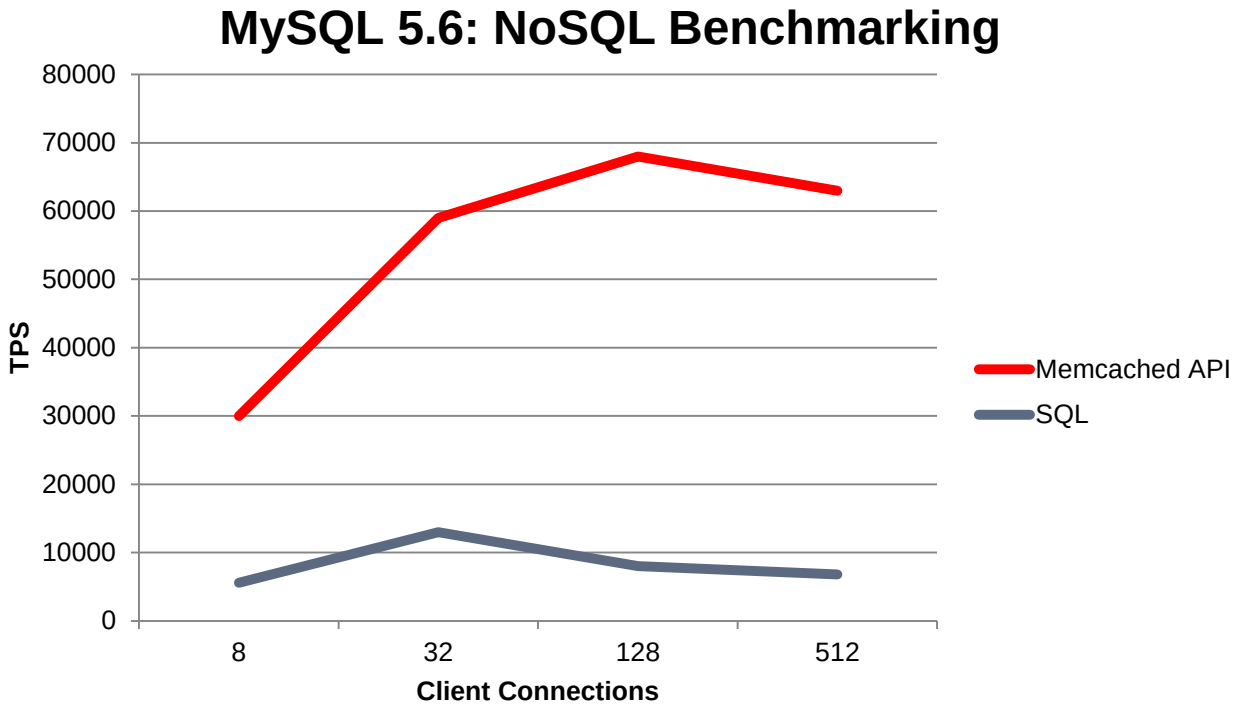

Up to 9x Higher "SET / INSERT" Throughput

MySQL Cluster: Multiple NoSQL Interfaces

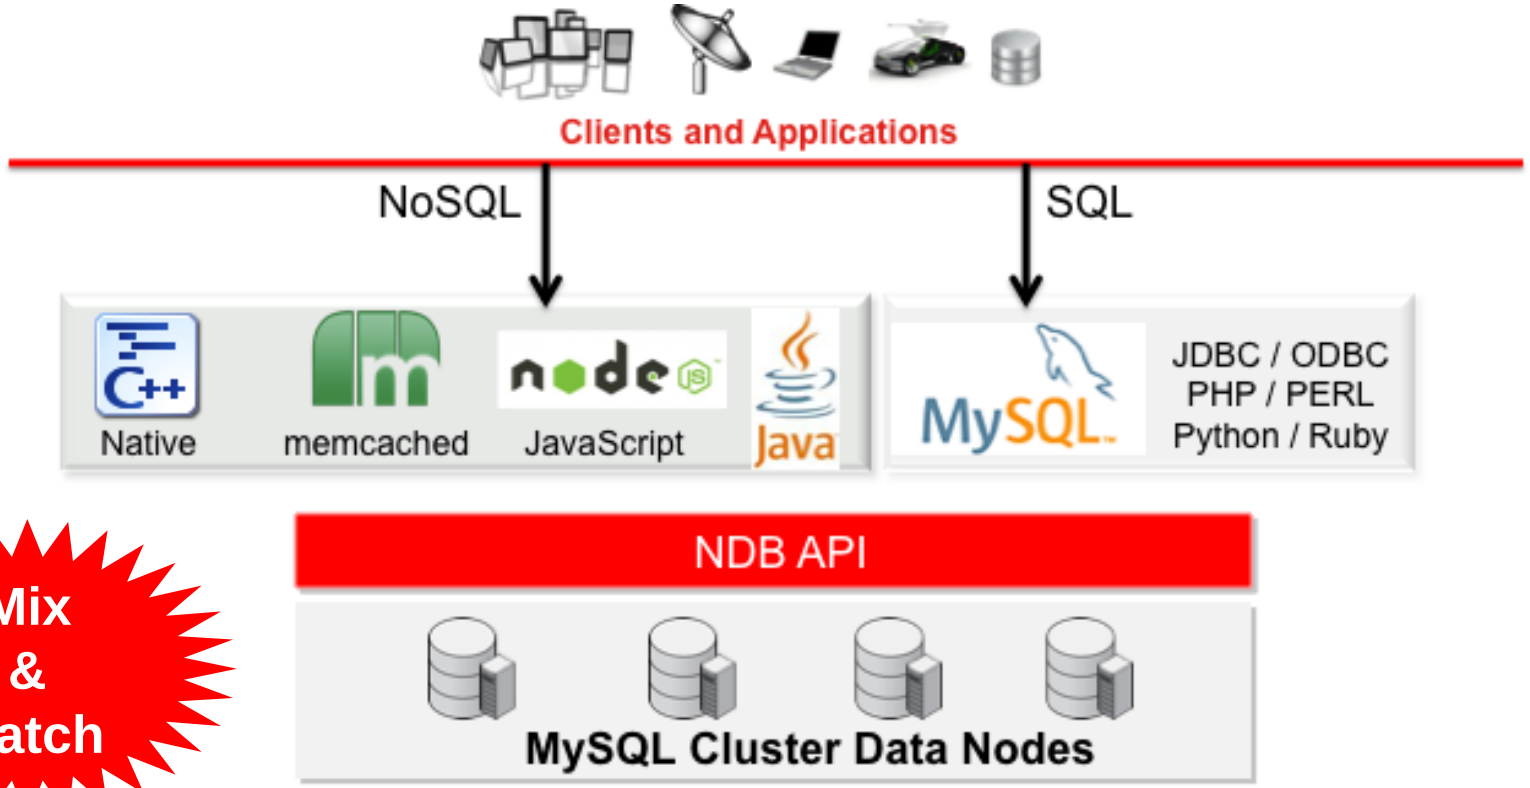

1.2 Billion UPDATEs per Minute

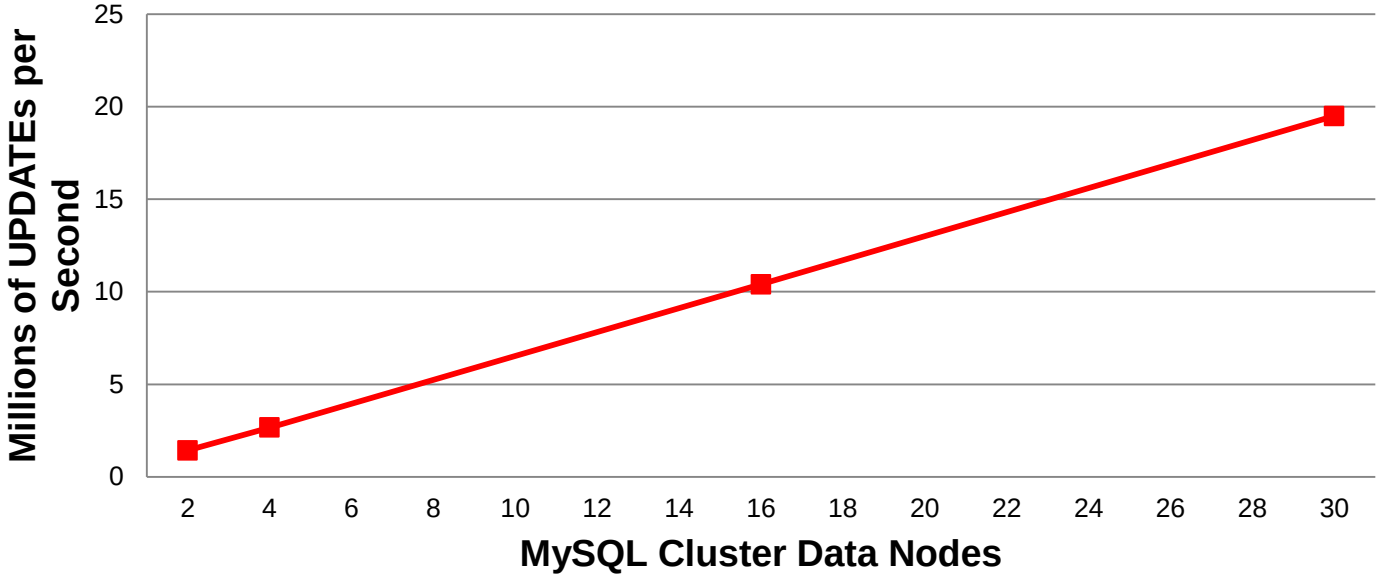

- NoSQL C++ API, flexaSynch benchmark
- 30 x Intel E5-2600 Intel Servers, 2 socket, 64GB
- ACID Transactions, with Synchronous Replication

MySQL in the Big Data Lifecycle

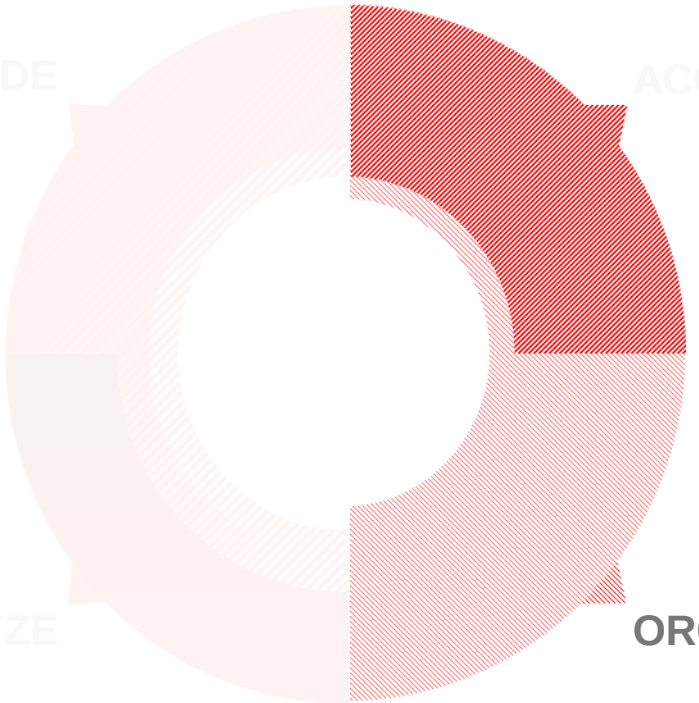

Import Data

Apache Sqoop
MySQL Hadoop Applier

Apache Sqoop

- Apache TLP, part of Hadoop project
  - Developed by Cloudera
- Bulk data import and export
  - Between Hadoop (HDFS) and external data stores
- JDBC Connector architecture
  - Supports plug-ins for specific functionality
- “Fast Path” Connector developed for MySQL

The Apache Software Foundation
http://www.apache.org/

Importing Data

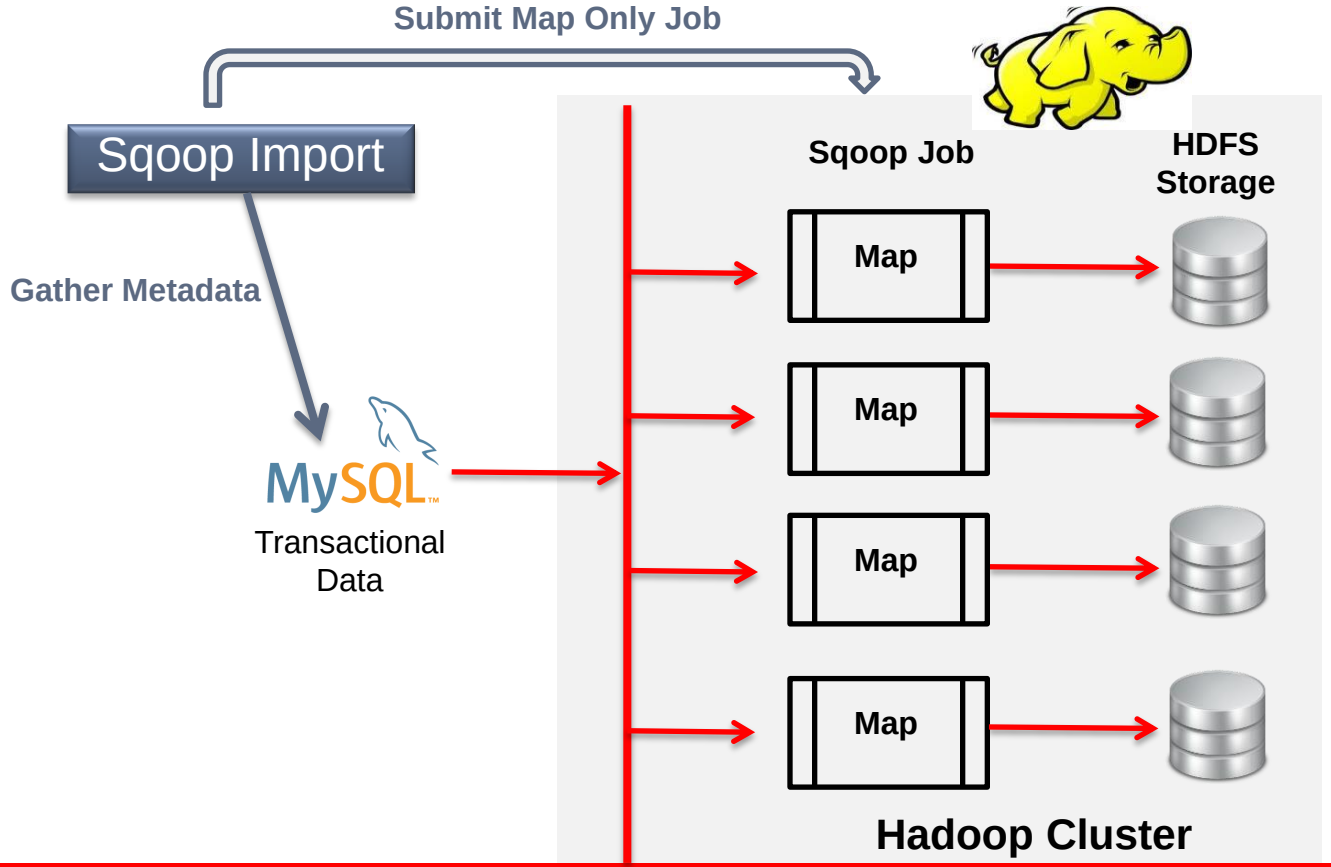

MySQL Applier for Hadoop

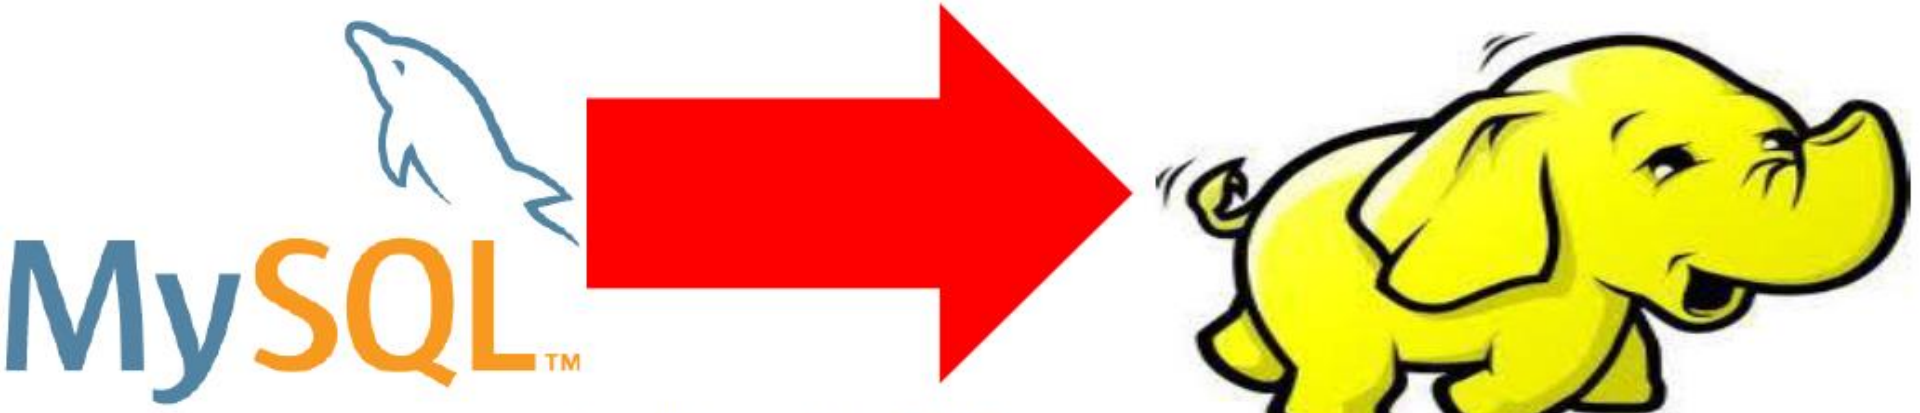

Real-Time Event Streaming
MySQL to HDFS

Hadoop Applier: Design

- Uses MySQL replication techniques for real time integration
- Binlog API uses Binary Log to rapidly fetch new data from a running server via the replication protocol
- MySQL Binlog comprised of events, each event represents a database change
- Hadoop Applier receives the events using the Binlog API, and writes the changes into a file in Hadoop Distributed File System
- Other tools in Hadoop Ecosystem,such as Apache Hive, can then consume this data

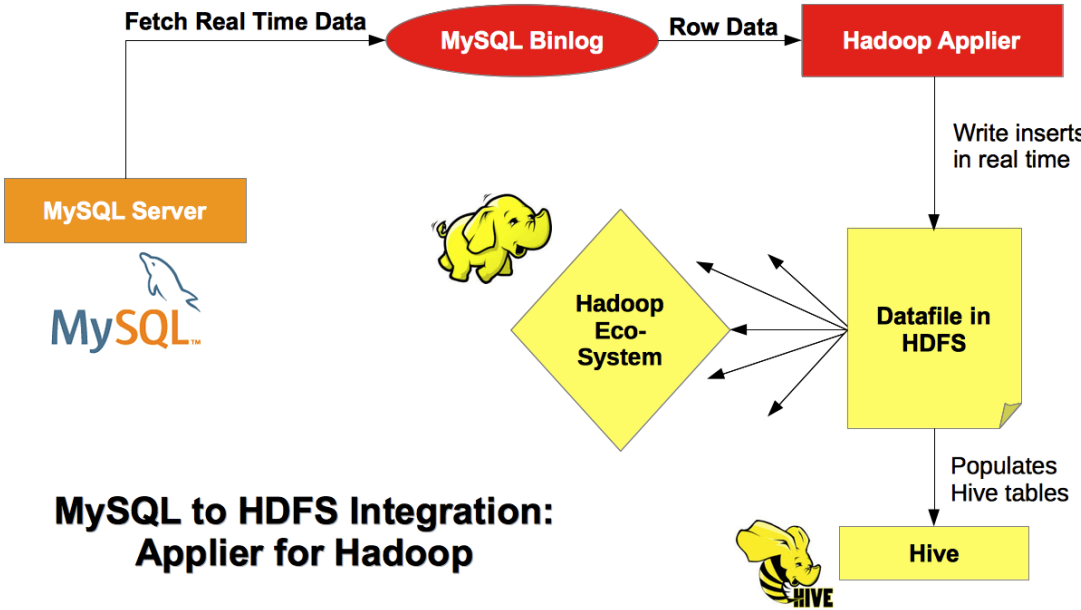

New Tool: MySQL Applier for Hadoop

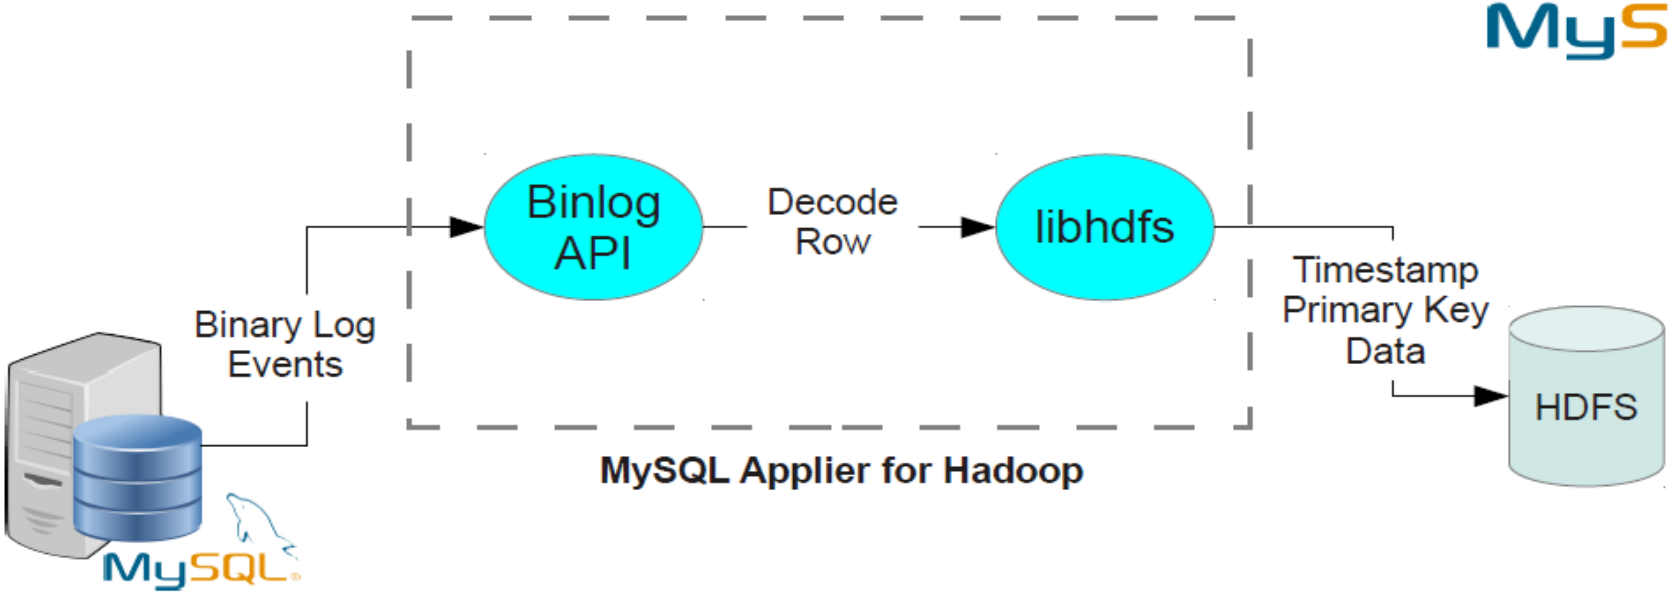

Hadoop Applier: Implementation

- Replicates rows inserted into a table in MySQL to Hadoop Distributed File System
- Uses an API provided by libhdfs, a C library to manipulate files in HDFS
- The library comes pre-compiled with Hadoop Distributions
- Connects to the MySQL master (or reads the binary log generated by MySQL) to:
  - Fetch the row insert events occurring on the master
  - Decode these events, extracting data inserted into each field of the row
  - Separate the data by the desired field delimiters and row delimiters
  - Use content handlers to get it in the format required
  - Append it to a text file in HDFS

Integration with HIVE

- Hive runs on top of Hadoop. Install HIVE on the hadoop master node
- Set the default datawarehouse directory same as the base directory into which Hadoop Applier writes
- Create similar schema's on both MySQL & Hive
- Timestamps are inserted as first field in HDFS files
- Data is stored in 'datafile1.txt' by default
- The working directory is base_dir/db_name.db/tb_name

SQL Query	Hive QL
CREATE TABLE t (i INT);	CREATE TABLE t (time_stamp INT, i INT) [ROW FORMAT DELIMITED] STORED AS TEXTFILE;

Mapping Between MySQL and HDFS Schema

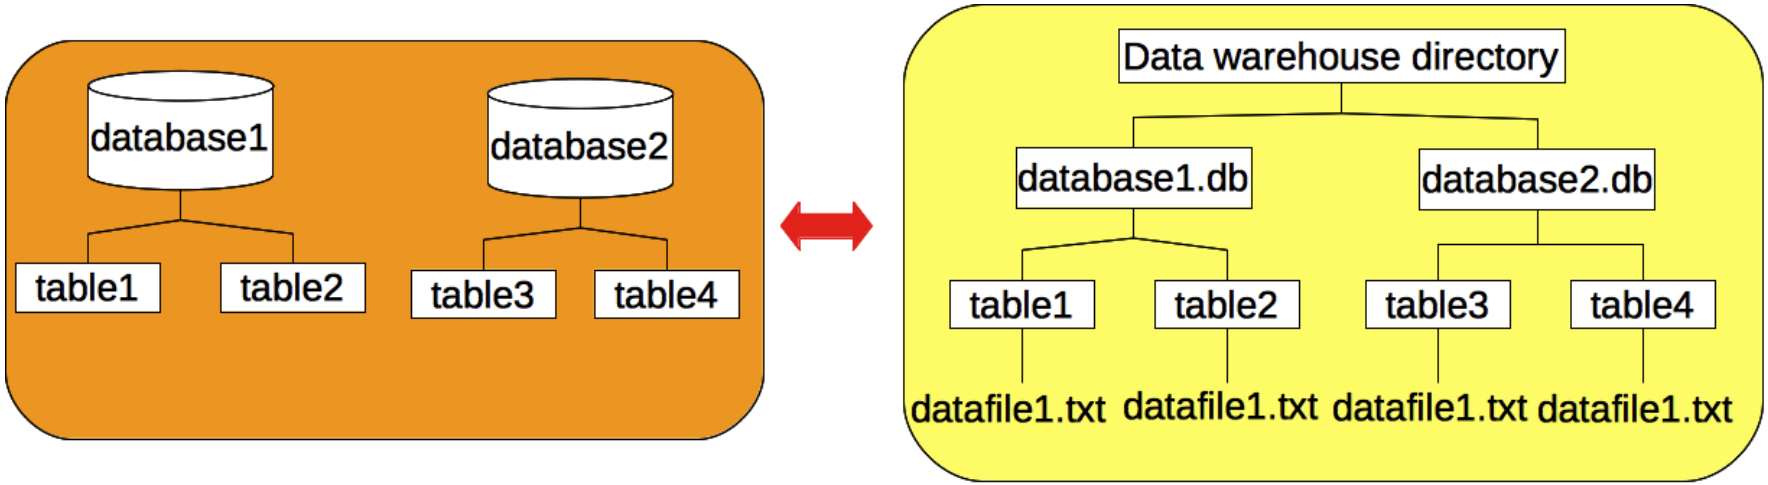

MySQL

table1

col1	col2	.....
data1	data2	....
data3	data4	....

table1/datafile1.txt

```
ts1,data1,data2,...
ts2,data3,data4,...
....
```

ts=timestamp

MySQL in the Big Data Lifecycle

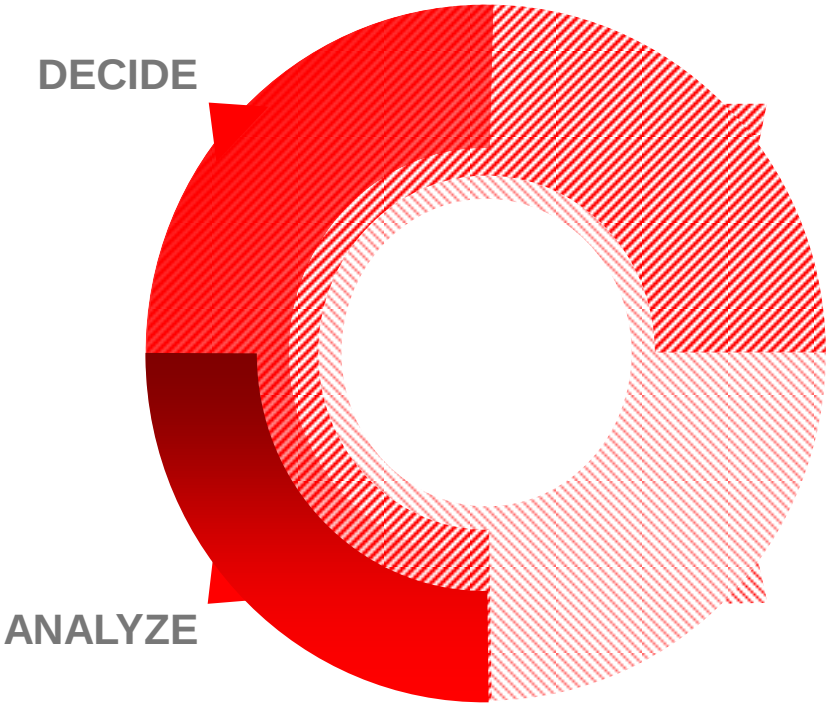

Analyze
Export Data
Decide

Analyze Big Data

Export Data

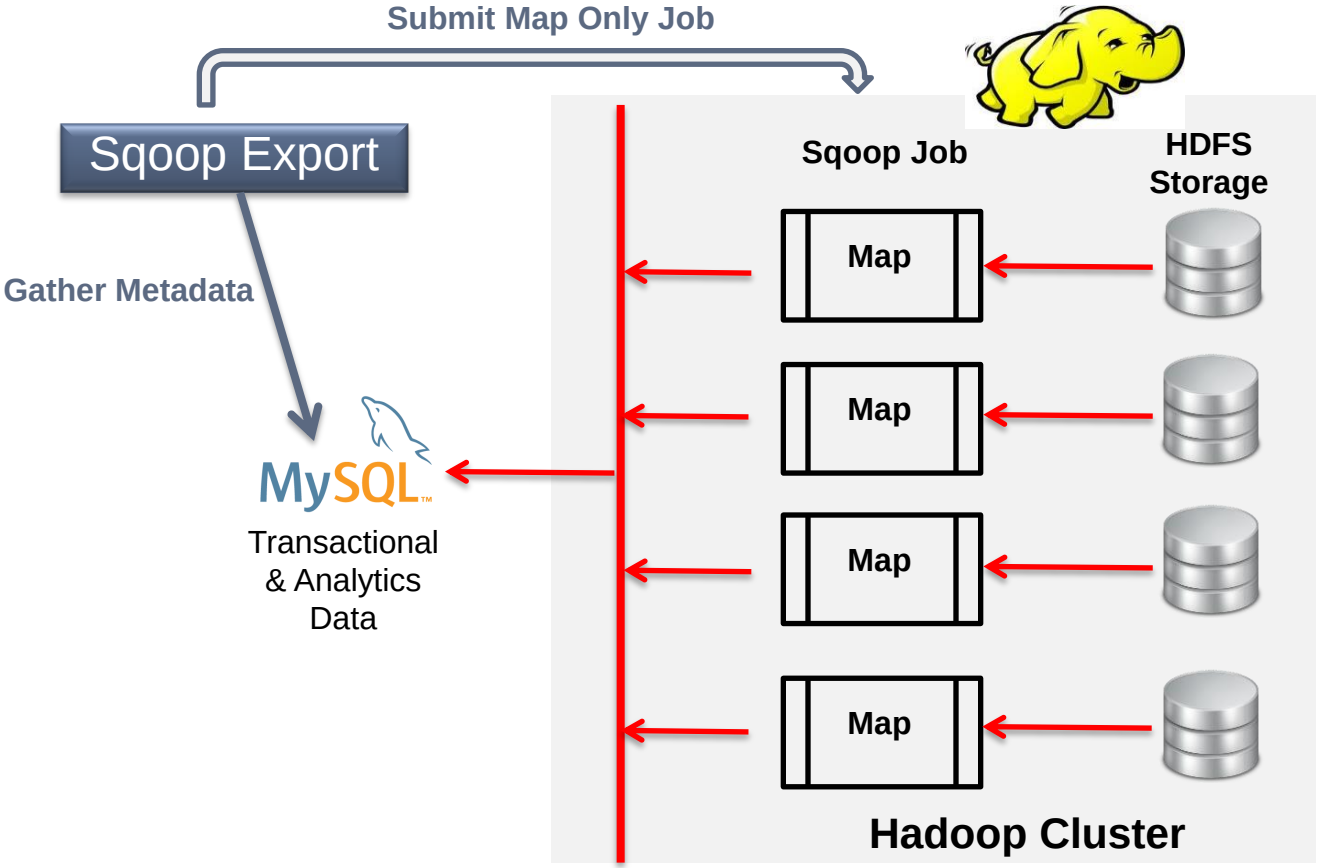

MySQL Reporting Database for BI

pentaho™
open source business intelligence™

MySQL™

talend*
*open data solutions

JASPERSOFT

MySQL Operational Database for Web

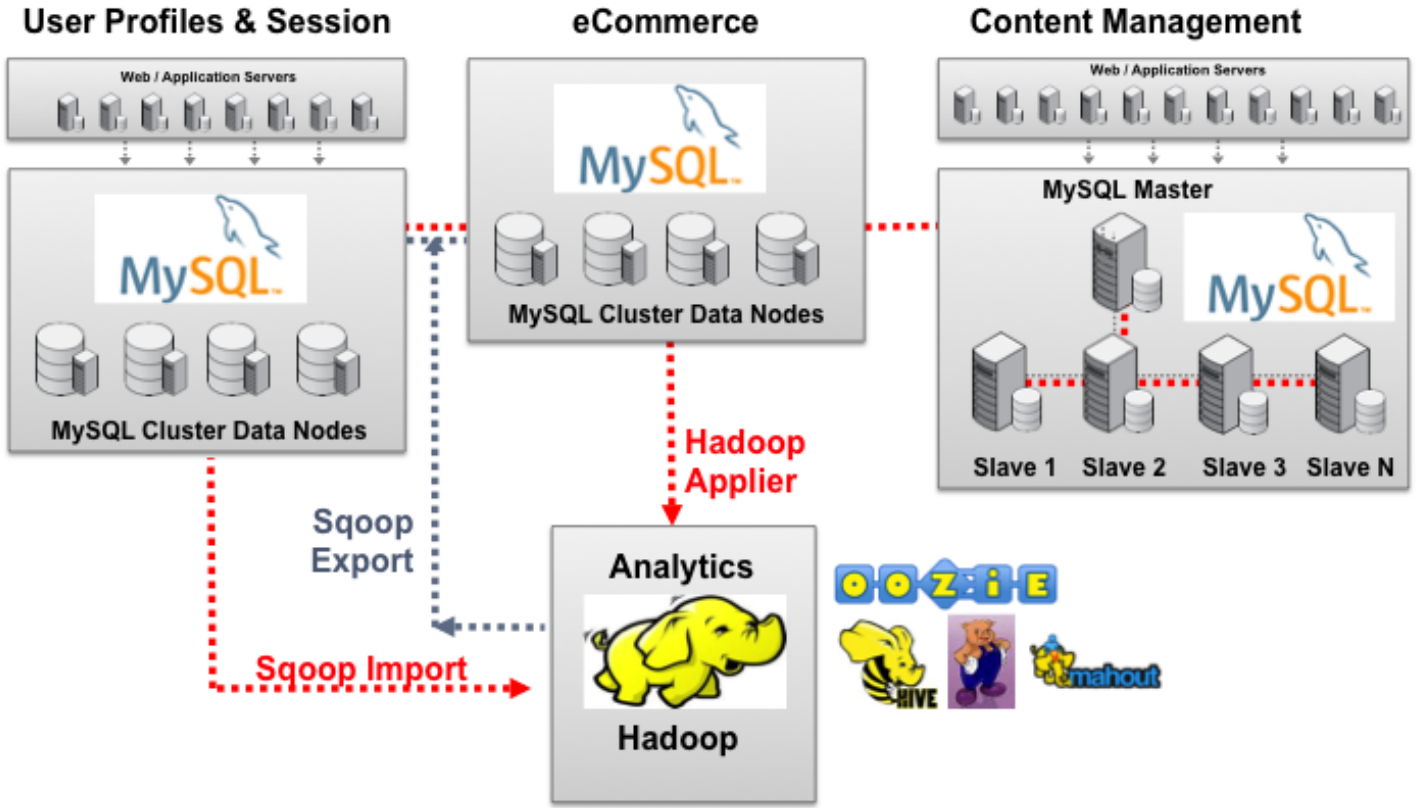

Data Analysis: MySQL Enterprise Edition

Highest Levels of Security, Performance and Availability

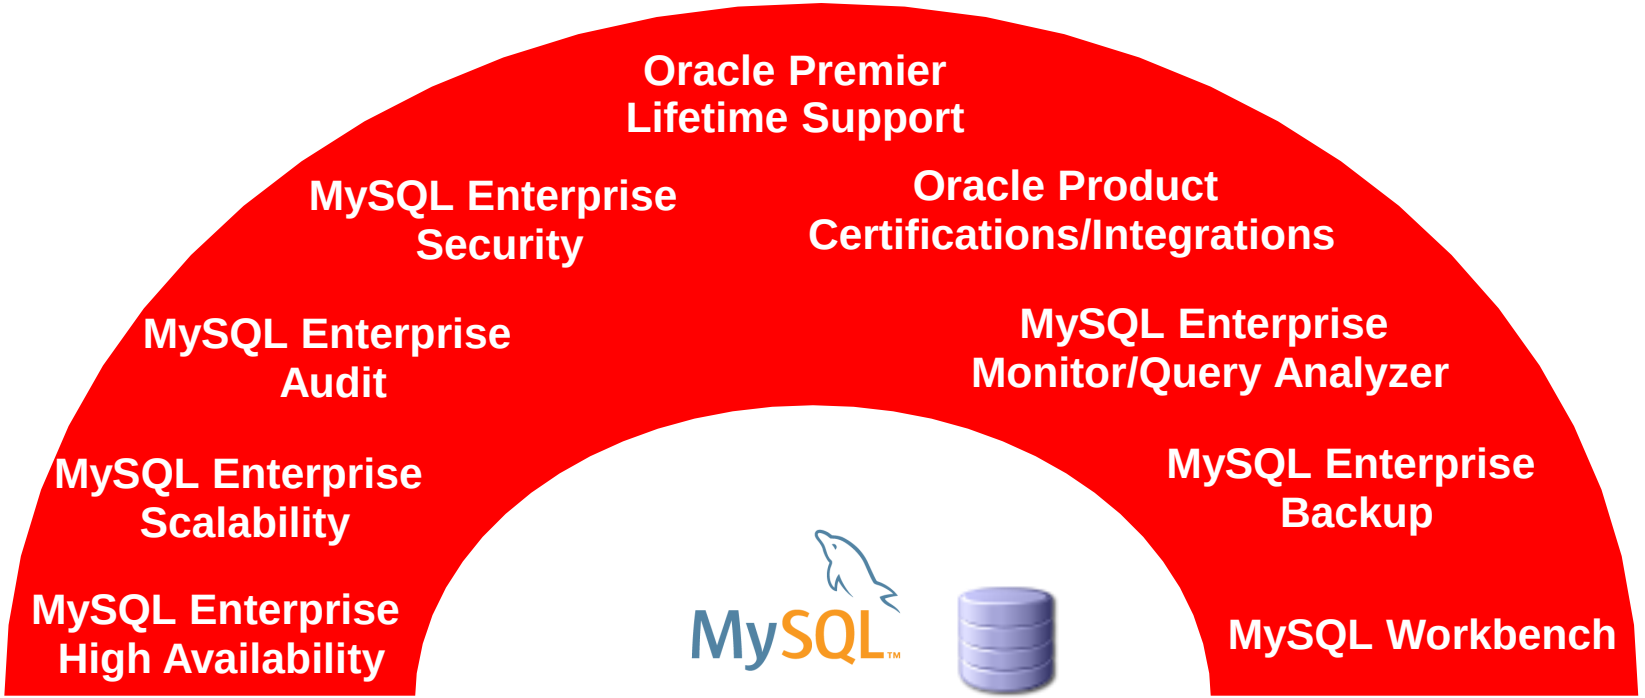

MySQL Enterprise Monitor with Query Analyzer

Tune Analytical Queries

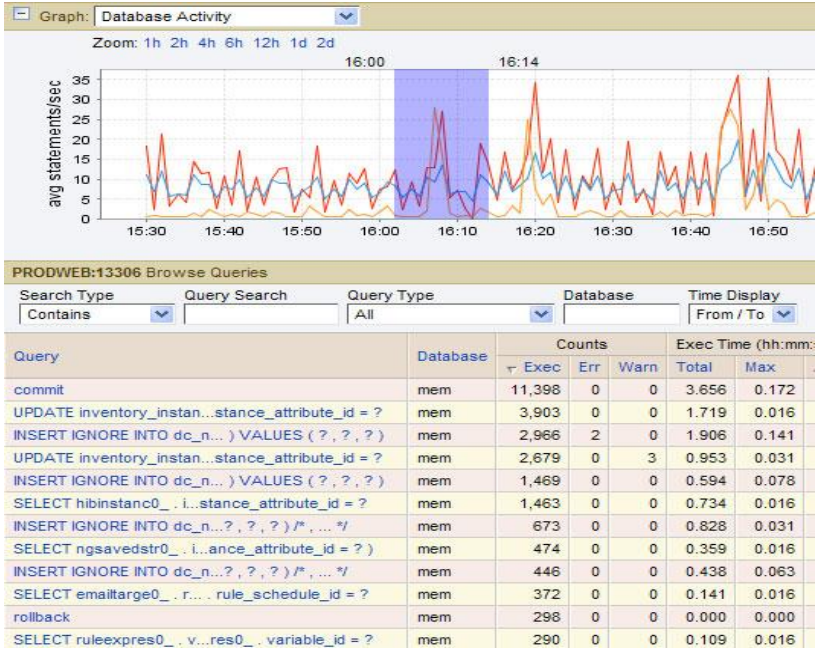

Query	Database	Counts			Exec Time (hh:mm:	
		Exec	Err	Warn	Total	Max
commit	mem	11,398	0	0	3.656	0.172
UPDATE inventory_instan...stance_attribute_id = ?	mem	3,903	0	0	1.719	0.016
INSERT IGNORE INTO dc_n...) VALUES (? , ? , ?)	mem	2,966	2	0	1.906	0.141
UPDATE inventory_instan...stance_attribute_id = ?	mem	2,679	0	3	0.953	0.031
INSERT IGNORE INTO dc_n...) VALUES (? , ? , ?)	mem	1,469	0	0	0.594	0.078
SELECT hibinstanc0_ . i...stance_attribute_id = ?	mem	1,463	0	0	0.734	0.016
INSERT IGNORE INTO dc_n...? , ? , ?) /* , ... */	mem	673	0	0	0.828	0.031
SELECT ngsavedstr0_ . i...ance_attribute_id = ?)	mem	474	0	0	0.359	0.016
INSERT IGNORE INTO dc_n...? , ? , ?) /* , ... */	mem	446	0	0	0.438	0.063
SELECT emailtarge0_ . r... . rule_schedule_id = ?	mem	372	0	0	0.141	0.016
rollback	mem	298	0	0	0.000	0.000
SELECT ruleexpres0_ . v...res0_ . variable_id = ?	mem	290	0	0	0.109	0.016

Enhance DevOps Agility

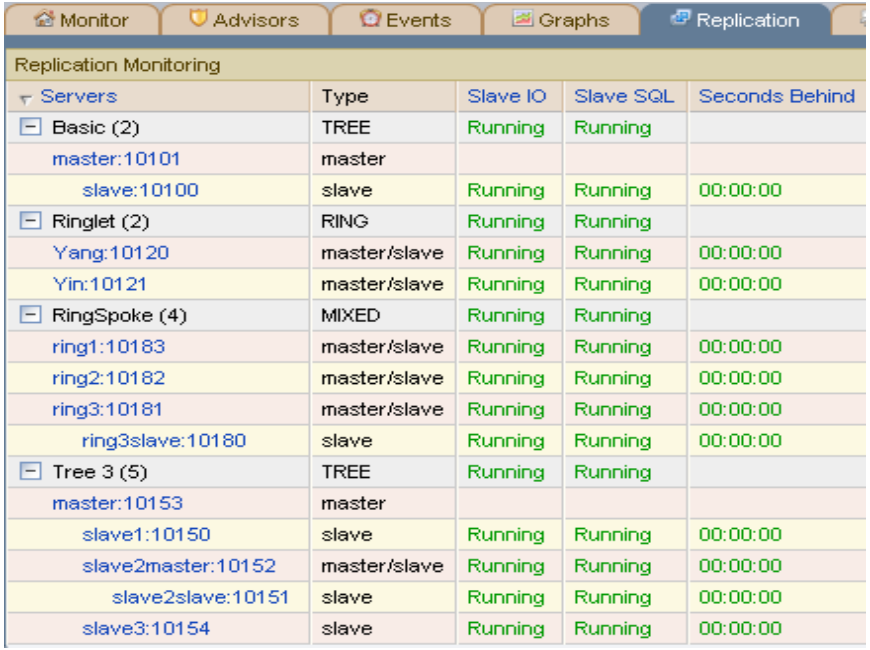

Servers	Type	Slave IO	Slave SQL	Seconds Behind
Basic (2)	TREE	Running	Running	
master:10101	master			
slave:10100	slave	Running	Running	00:00:00
Ringlet (2)	RING	Running	Running	
Yang:10120	master/slave	Running	Running	00:00:00
Yin:10121	master/slave	Running	Running	00:00:00
RingSpoke (4)	MIXED	Running	Running	
ring1:10183	master/slave	Running	Running	00:00:00
ring2:10182	master/slave	Running	Running	00:00:00
ring3:10181	master/slave	Running	Running	00:00:00
ring3slave:10180	slave	Running	Running	00:00:00
Tree 3 (5)	TREE	Running	Running	
master:10153	master			
slave1:10150	slave	Running	Running	00:00:00
slave2master:10152	master/slave	Running	Running	00:00:00
slave2slave:10151	slave	Running	Running	00:00:00
slave3:10154	slave	Running	Running	00:00:00

Scaling, Security and Data Protection

- MySQL Enterprise Scalability
- MySQL Enterprise Security
- MySQL Enterprise Audit
- MySQL Enterprise Backup

MySQL Enterprise Backup

- Online Backup for InnoDB
- Full, Incremental, Partial Backups (scriptable interface)
- Compression
- Point in Time, Full, Partial Recovery options
- Metadata on status, progress, history
- Unlimited Database Size
- Cross-Platform
  - Windows, Linux, Unix
- Certified with Oracle Secure Backup

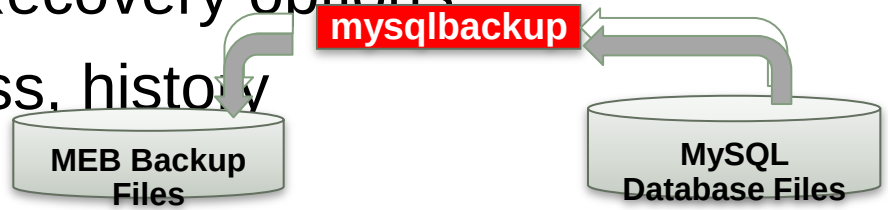

ORACLE

*Ensures **quick, online backup and recovery** of your MySQL apps.*

MySQL Enterprise Security

MySQL External Authentication

- PAM (Pluggable Authentication Modules)
  - Access external authentication methods
  - Standard interface (Unix, LDAP, others)
  - proxied and non-proxied users
- Windows
  - Access native Windows services
  - Authenticate users already logged into Windows (Windows Active Directory)
- Pluggable Authentication API

Integrates MySQL with existing security infrastructures and SOPs

5.5 MySQL Enterprise Scalability

MySQL Thread Pool

- MySQL default thread-handling – excellent performance, can limit scalability as connections grow
- MySQL Thread Pool improves sustained

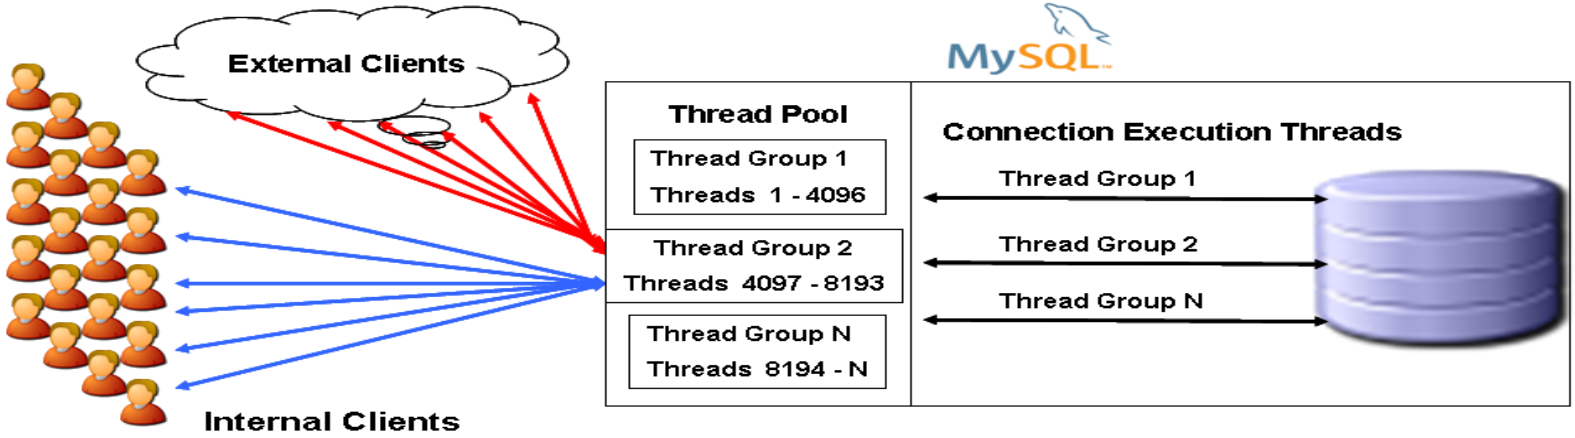

Thread Pool

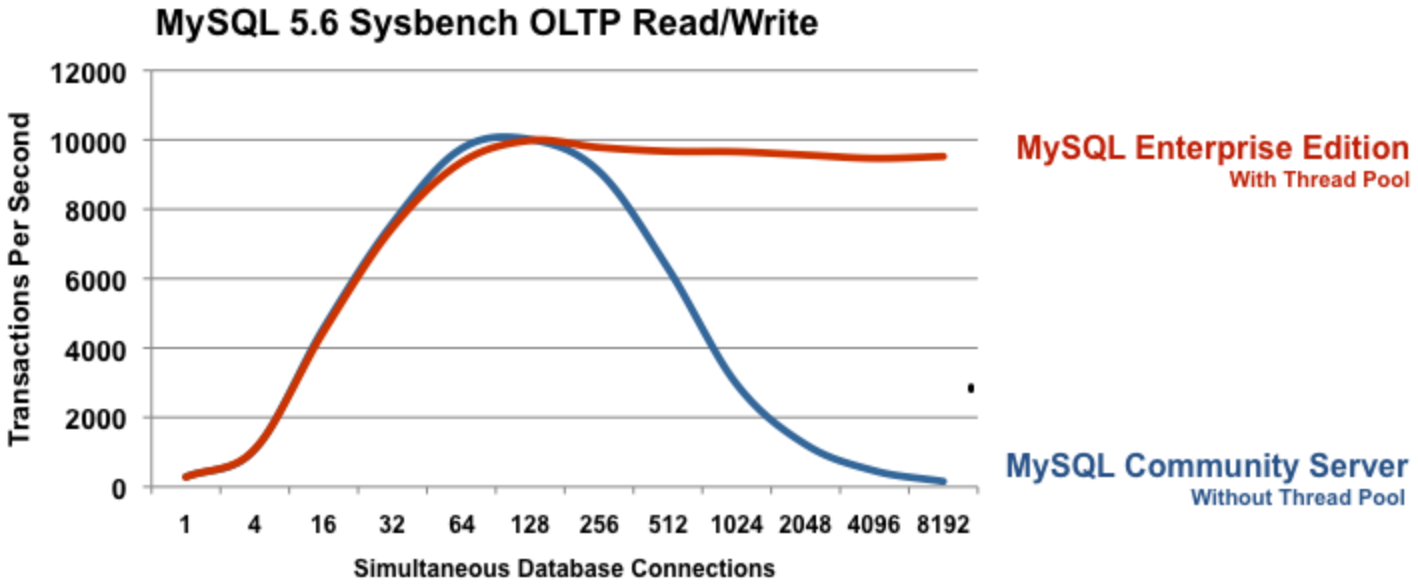
MySQL 5.6 Sysbench OLTP Read/Write
Transactions Per Second
12000
10000
8000
6000
4000
2000
0
1
4
16
32
64
128
256
512
1024
2048
4096
8192
Simultaneous Database Connections
MySQL Enterprise Edition
With Thread Pool
MySQL Community Server
Without Thread Pool

MySQL Enterprise Audit

Policy-based Auditing for MySQL Applications

- Out-of-the-box logging of connections, logins, query activity across all or specific MySQL servers
- User defined policies, filtering and log rotation
- Dynamically enabled, disabled: no server restart
- XML-based audit stream per Oracle audit specification
- Easily implemented via MySQL 5.5 Audit API
- MySQL 5.5.28 and higher

Get it here: support.oracle.com and edelivery.oracle.com

Adds regulatory compliance to MySQL applications

Oracle Premier Support for MySQL

Rely on The Experts - Get Unique Benefits

- Straight from the Source
- Largest Team of MySQL Experts
- Backed by MySQL Developers
- Forward Compatible Hot Fixes
- MySQL Maintenance Releases
- MySQL Support in 29 Languages
- 24/7/365
- Unlimited Incidents
- Knowledge Base
- MySQL Consultative Support

Only From Oracle

"The MySQL support service has been essential in helping us with troubleshooting and providing recommendations for the production cluster, Thanks."

-- Carlos Morales – Playfulplay.com

ORACLE

Summary

- MySQL + Hadoop: widely deployed solution
- “Best of both worlds” SQL + NoSQL Access
- Tools and expertise to support you
- End to end Oracle solutions for Big Data

Integrate for Insight

Next Steps

Download the Guide

http://www.mysql.com/why-mysql/white-papers/mysql_wp_hadoop.php

Try Out MySQL 5.6

http://www.mysql.com/downloads/mysql/

Engage MySQL Consulting

http://www.mysql.com/consulting/

ORACLE

Thank you!

Hardware and Software

ORACLE

Engineered to Work Together


ORACLE®