What's New in MySQL and MongoDB Ecosystem

Year 2017

Peter Zaitsev
CEO
Percona University, Ghent
June 22nd, 2017

PERCONA

In This Presentation

Few Words about Percona

Few Words about Percona University Program

Most Interesting Developments in MySQL and MongoDB ecosystems

Thank you – Co.Station Gent

PERCONA

Few Words about Percona

Percona's Purpose

To Champion Unbiased Open Source Database Solutions

We Do

Support, Managed Services for MySQL and MongoDB

Also Consulting and Training

Helping companies to migrate to Open Source Database

Develop Open Source Software

Solutions to maximize your success

Broad Software Ecosystem Support

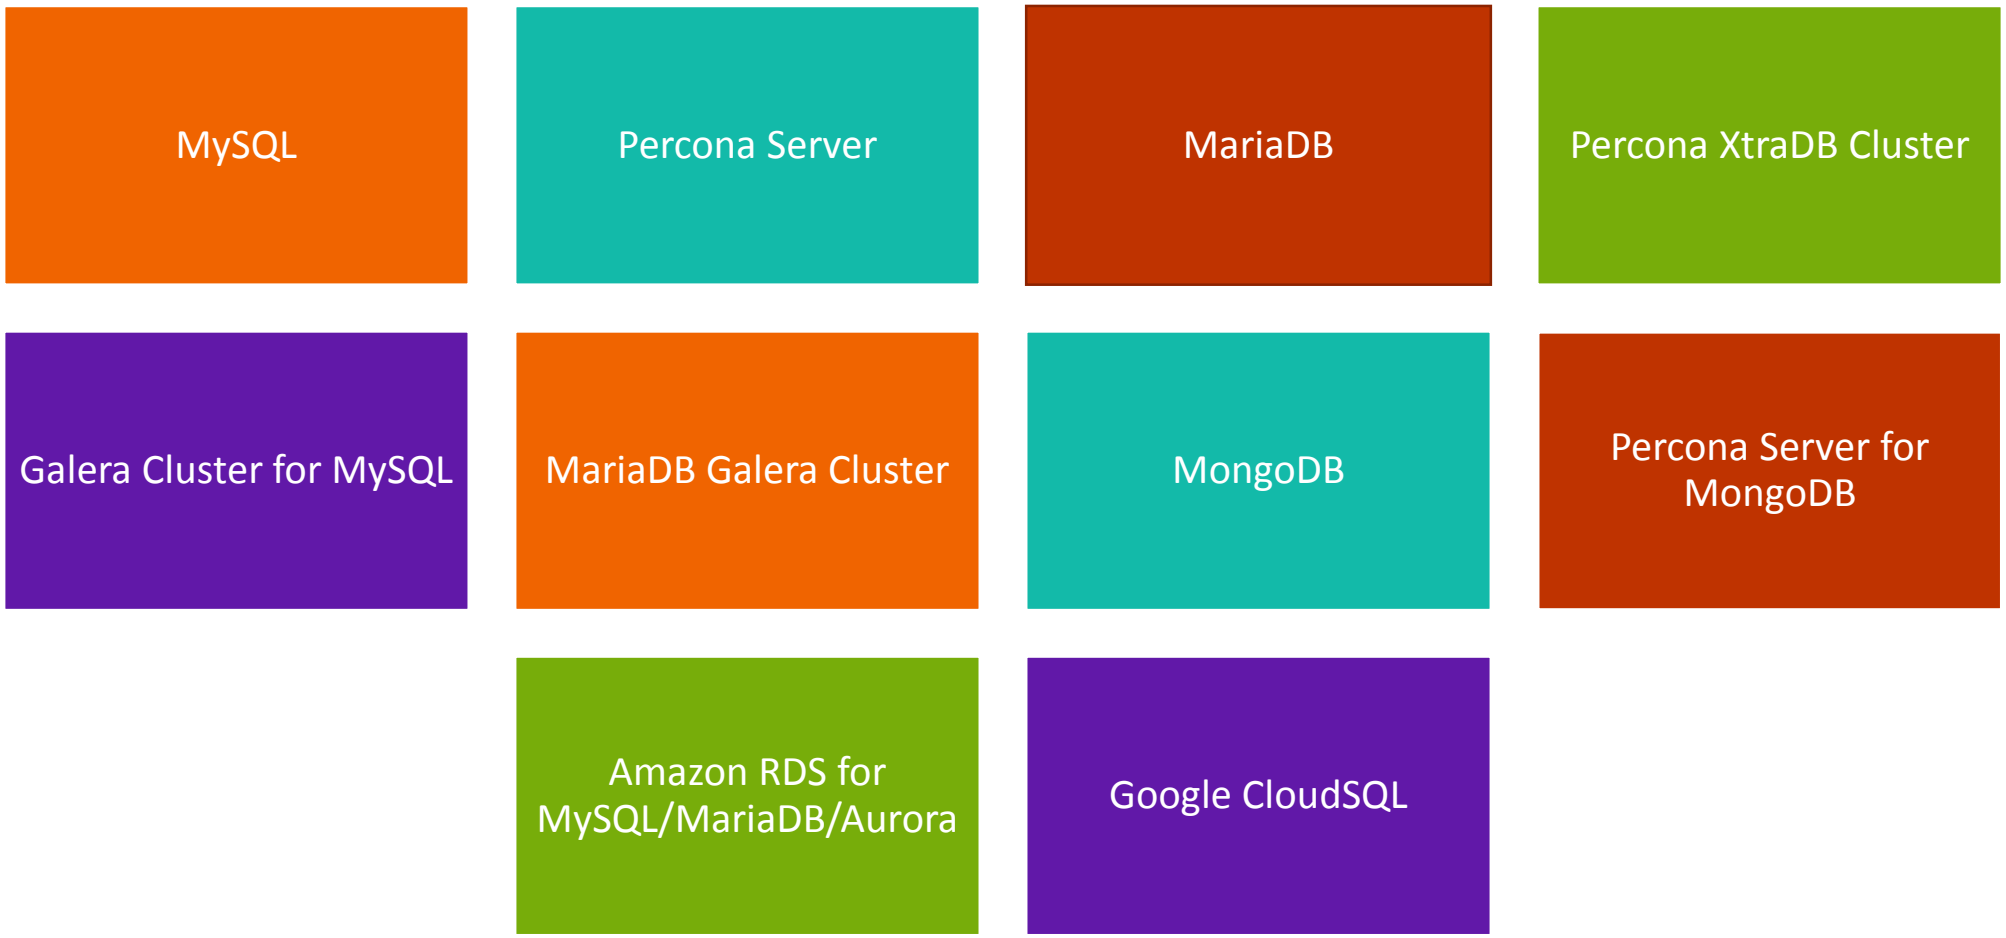

100% Free and Open Source Software

Percona Server for MySQL

Percona Server for MongoDB

Percona XtraDB Cluster

Percona Xtrabackup

Percona Toolkit

Percona Monitoring and Management

PERCONA

Percona University

Educational Technical Presentations

Multiple Locations in the World

Partnering with Local Companies

Affordable to Attend (Free)

Percona University – What to Expect

- Several presentations on different topics
- Feel free to only attend those you're interested in
- Keep it Interactive! Ask Questions
- Breaks
- Prize Give away in the end

Whats new in MySQL and MongoDB

Top Highlights

Innovations worth Noticing

MySQL 5.7	MySQL 8	MariaDB 10.1	Amazon Aurora
Percona XtraDB Cluster 5.7	Percona Monitoring and Management	MyRocks	ProxySQL
Orchestrator	Gh-ost	MongoDB 3.4	Percona Server for MongoDB 3.4

MySQL 5.7 – Current GA

MySQL 5.7 - Security

"Secure by Default"

Password validation (no weak passwords by default)

Automatic SSL certificate creation

Simple SSL setup

"root" user created with password by default

MySQL 5.7 Encryption

Can encrypt Innodb tables on disk

Only data is encrypted at this point

Innodb log files, binary log files are not encrypted

MySQL 8 does Innodb log file encryption

MySQL 5.7 - NoSQL

Native JSON data type

Can index fields in JSON documents

CRUD access through Protocol X

MySQL 5.7 - Replication

- Parallel Replication
- Multi-Source Replication
- Can enable GTID online
- MySQL Group Replication
- MySQL Innodb Cluster

MySQL 5.7 – Performance Schema

- Automatic configuration for Performance Schema
- Reduced overhead (especially memory overhead)
- Memory Usage Instrumentation
- Instrumentation of Storage Procedures
- Instrumentation of Transactions
- Sys_schema included for simple Performance Schema access

MySQL 5.7 - Performance

Further improved Multi-core scalability

Optimizations for Innodb Temporary Tables

New Compression for Innodb Tables

Optimizer Improvements

MySQL 5.7 Benchmarketing: Sysbench OLTP Read Write

1.5x Faster than MySQL 5.6
2.5x Faster than MySQL 5.5

17,000 TPS

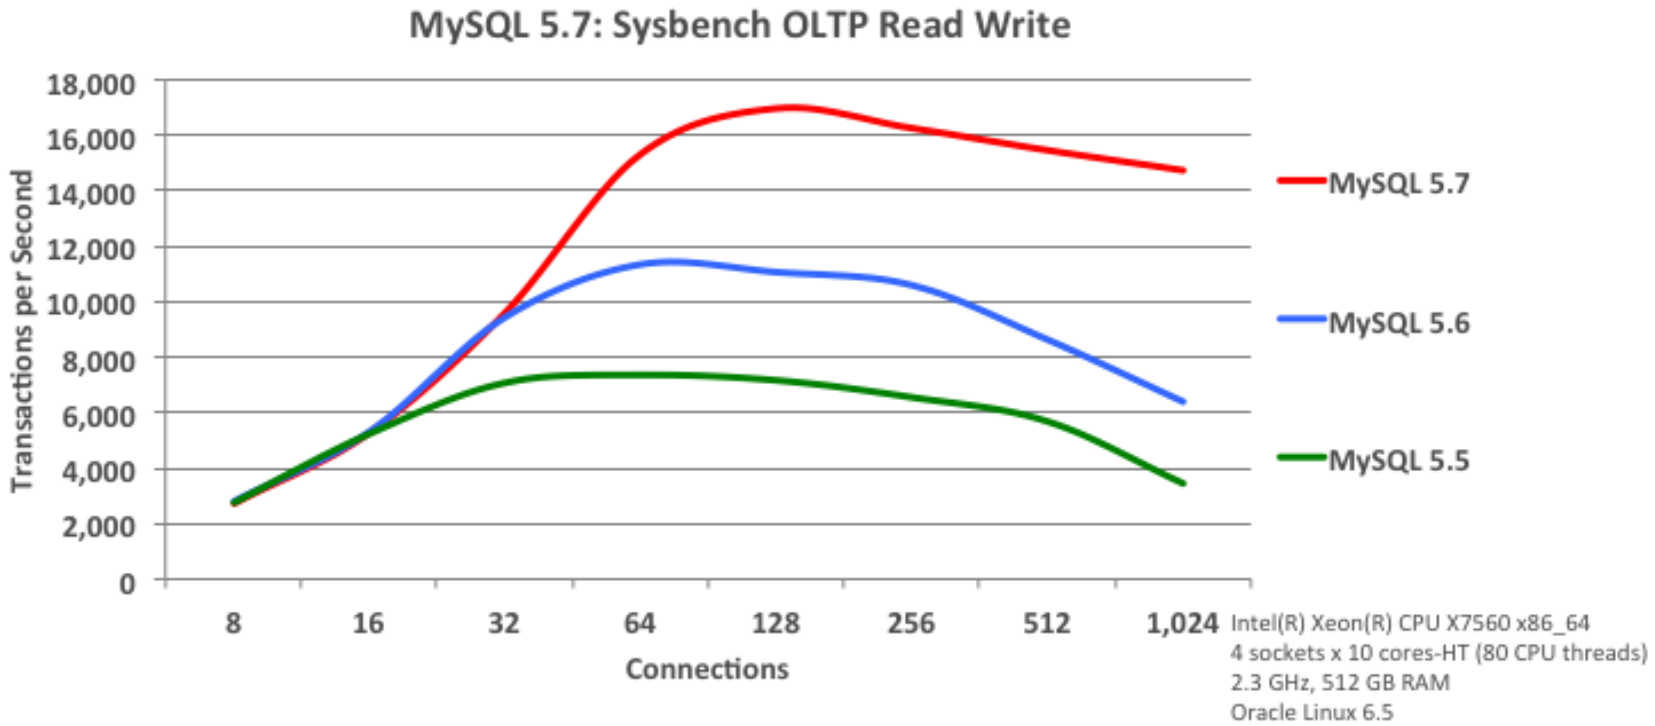

**Information from Oracle OpenWorld presentation by Geir Hoydalsvik*

PERCONA

Complete list of MySQL 5.7 Improvements

http://www.thecompletelistoffeatures.com/

MySQL 8 – Currently in Development

MySQL 8

- What was previously referred as MySQL 5.8
- Native Data Dictionary (no more .frm files)
- Roles
- Further Performance Schema Improvements
- Common Table Expressions (CTE) and Windowing Functions
- Replication Performance Improvements
- Invisible Indexes
- More to come

MariaDB 10.2 - Current GA

About MariaDB

"fork" by MySQL Founder Michael Widenious,

Replaced MySQL in many Linux Distributions

Has number of interesting features MySQL does not

More and More Becomes separate Database

Not every MySQL feature ported to MariaDB

Differences in Replication, Optimizer, JSON, Protocol X, GIS, Enryption

New in MariaDB 10.2

- Recursive Common Table Expresions (CTE) WITH
- CHECK CONSTRAINT support
- DEFAULT expressions
- DEFAULT value for TEXT/BLOB
- Alpha (Experimental) MyRocks Storage Engine
- JSON, GeoJSON improvements

Amazon Aurora

Amazon Aurora

- The fastest Growing Database Technology at AWS
- High End of Amazon RDS MySQL
- Integration with Amazon Cloud Storage for Improved Performance and Replication
- Automatic Replication and Cluster Recovery
- Improved Query Cache
- Not always faster than MySQL on EC2

Percona Server

Percona Server 5.7

- Full compatibility with MySQL Community Edition
- Many equivalents to MySQL Enterprise features
- Improved Innodb Performance on heavy load
- New design of DoubleWrite Buffer for better IO performance
- Column Compression (with custom dictionary support)
- Improved TokuDB Storage Engine

Percona XtraDB Cluster 5.7

Percona XtraDB Cluster 5.7

HA Solution based on Percona Server 5.7 and Galera Library

High Availability for MySQL without pains of Async replication

Automated Node Provisioning and Self Healing

Works great in the Cloud and with Containers

Integration with ProxySQL for traffic management

Improved Ease of Use, Security and Performance

PXC 5.7 Performance Improvements

http://bit.ly/2qGCr0T and http://bit.ly/2pzvAIW

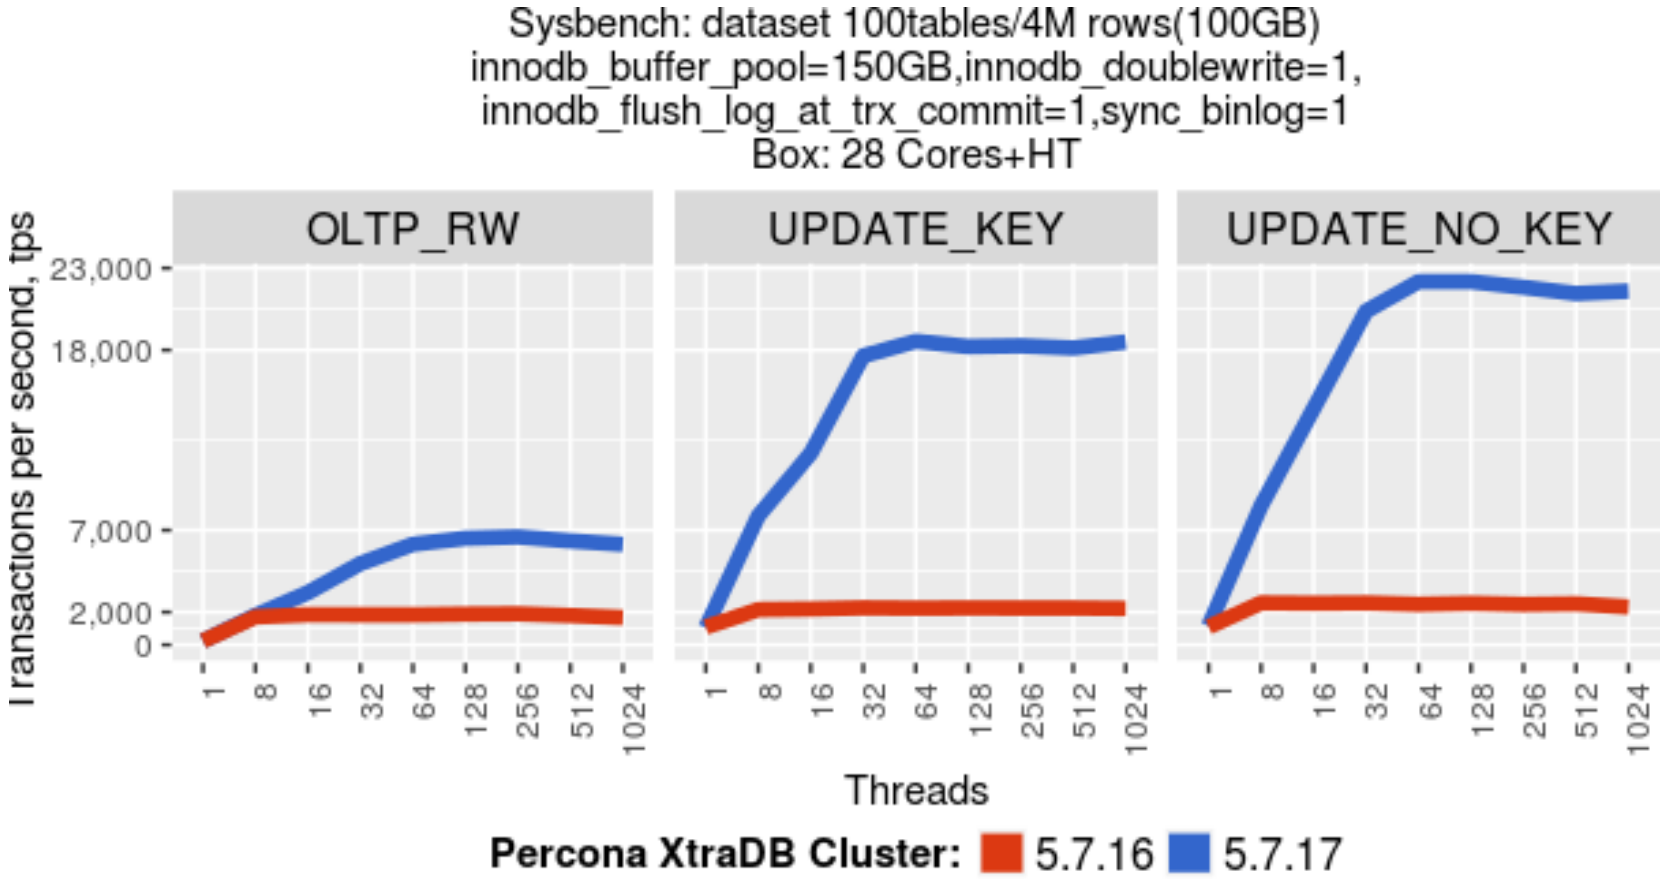

PERCONA

MyRocks

MyRocks

MySQL Storage engine based on MyRocks

Used by Facebook for better efficiency and performance

Uses LSM trees as underlying data structure

Write Optimized Engine

Experimental integration available with Percona Server and MariaDB

RocksDB Efficiency

Mark Callaghan: http://bit.ly/2epDJqD

Small server: Linkbench, IO-bound

MyRocks: best throughput & QoS, most efficient
MongoRocks: better than WiredTiger

	TPS	iostat r/t	iostat wKB/t	CPU usecs/t	Size (GB)	p99 update
MongoRocks+zlib	1087	1.07	4.42	24656	23	2
WiredTiger+zlib	429	1.24	17.98	153763	33	22
MyRocks+zlib	**2246**	**0.67**	**1.27**	12688	**18**	**1**
InnoDB	1860	0.82	10.62	**7991**	63	14
InnoDB+zlib	1855	**0.67**	8.60	10431	40	8

Percona Monitoring and Management

Percona Monitoring and Management

100% Free and Open Source

Comprehensive Database focused Monitoring

Supports MySQL and MongoDB (and variants)

Easy to Install and Use

Version 1.x focuses on Trending and Query Analyses

Management Features to come

Why did we created PMM

No good database focused monitoring solution

Existing solutions Proprietary or Cloud Only

Do it yourself is possible but hard

Want everyone to be able to run monitoring in every environment

To fix problems before they appear

What Queries are causing the load ?

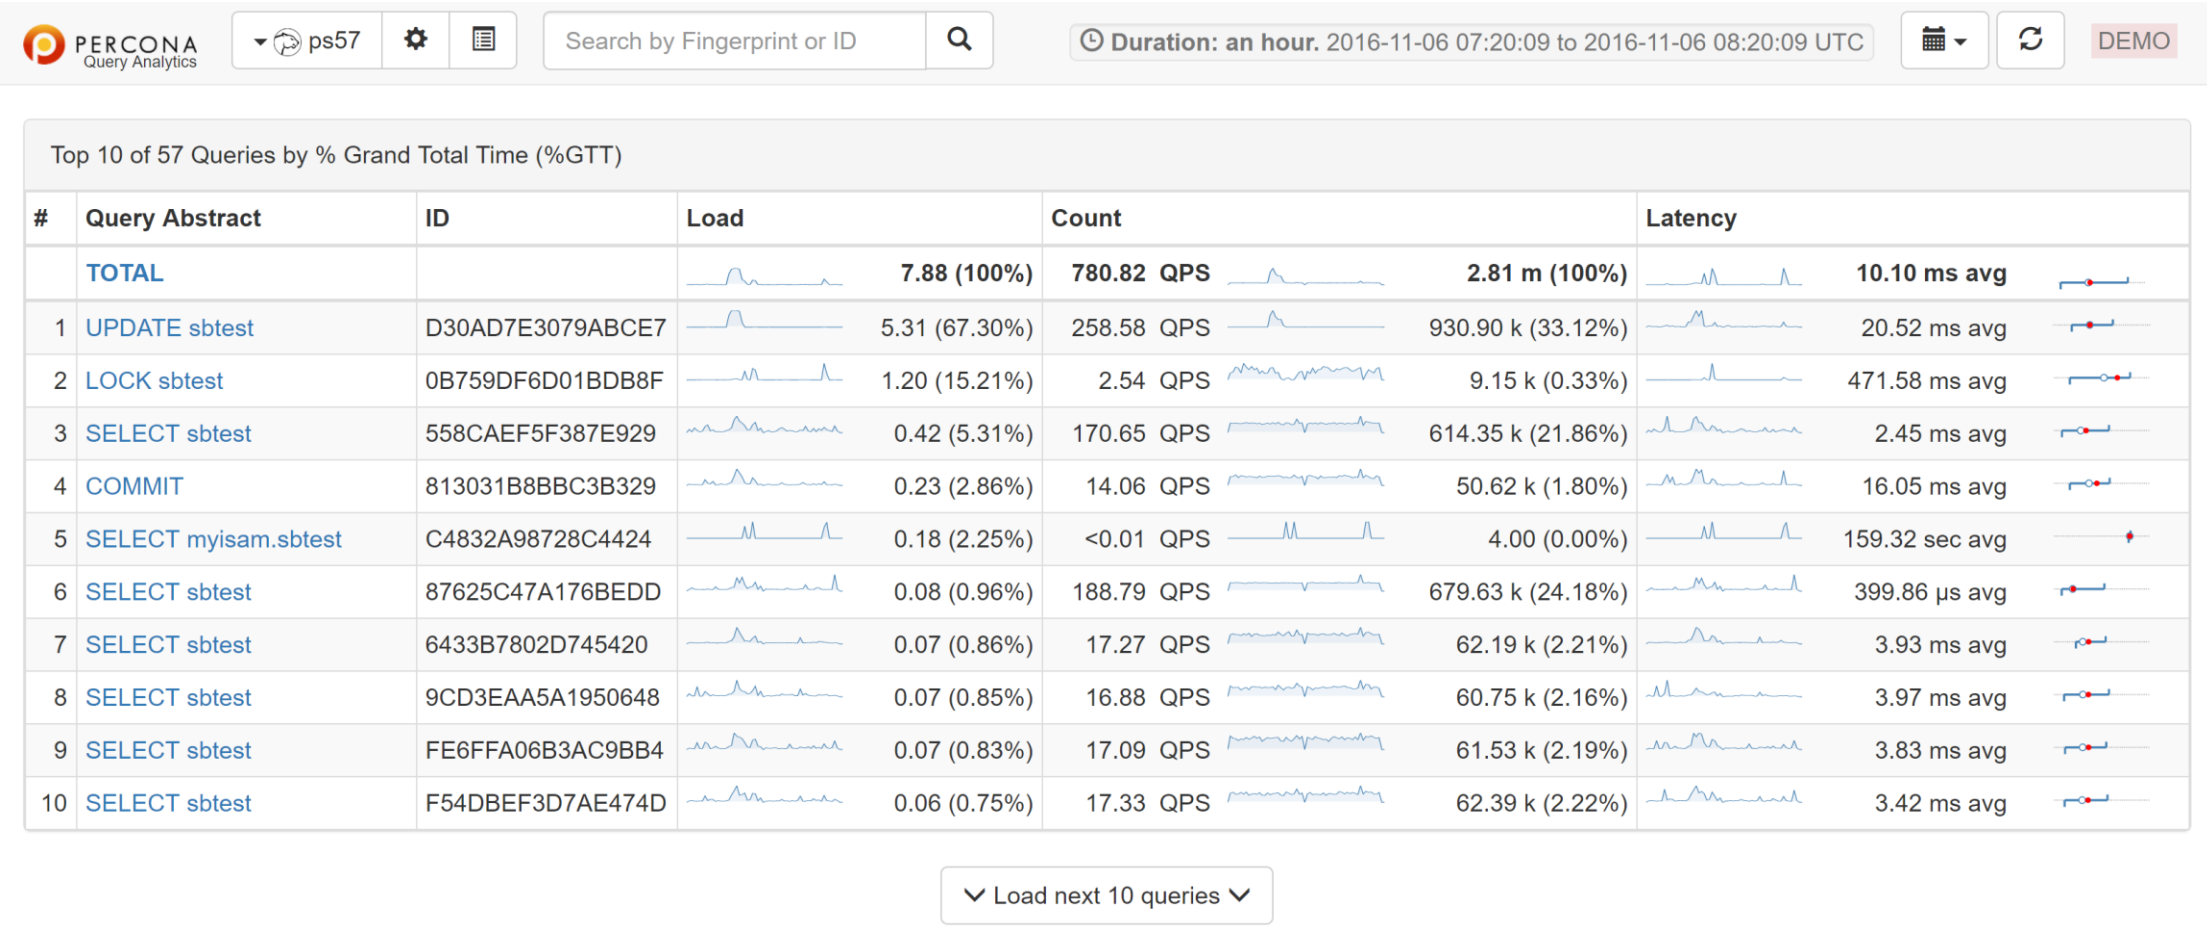

Why are they causing this load ?

UPDATE sbtest D30AD7E3079ABCE7

Selected query class: 930.90 k Queries (258.58 QPS, 67.30%, 5.31 Load) | Total: 2.81 m Queries (780.82 QPS, 100.00%, 7.88 Load)

Metrics	Rate/Sec	Sum	Per Query Stats
Query Count	259.12 (per sec)	932.85 k 32.40% of total	
Query Time	5.31 load (67.16%)	19101.45 sec 67.16% of total	5.25 ms avg
Lock Time	1.38 (avg load)	4965.60 sec 51.50% of total 15.13% of query time	793.79 µs avg
Innodb Row Lock Wait	<0.01 (avg load)	28.14 sec 42.03% of total 1.44% of query time	75.79 µs avg
Innodb IO Read Wait	<0.01 (avg load)	35.85 sec 2.47% of total 10.93% of query time	573.60 µs avg
Innodb Read Ops	2.48 (per sec)	8.91 k 1.86% of total	0.00 avg
Innodb Read Bytes	39.61 KB (per sec)	139.25 MB 1.86% of total 16.00 KB avg io size	3.50 KB avg
Innodb Distinct Pages	-	-	6.03 avg
Bytes Sent	13.18 KB (per sec)	46.35 MB 1.72% of total	52.00 Bytes avg
Rows Examined	258.17 (per sec)	929.43 k 0.64% of total 0.00 per row sent	0.88 avg

How to fix them

EXPLAIN

Database: innodb EXPLAIN

Id	SelectType	Table	Partitions	CreateTable	Type	PossibleKeys	Key	KeyLen	Ref	Rows	Extra
1	SIMPLE	sbtest1			const	PRIMARY	PRIMARY	4	const	1	

CREATE STATUS

```
CREATE TABLE `sbtest1` (
  `id` int(10) unsigned NOT NULL AUTO_INCREMENT,
  `k` int(10) unsigned NOT NULL DEFAULT '0',
  `c` char(120) NOT NULL DEFAULT '',
  `pad` char(60) NOT NULL DEFAULT '',
  PRIMARY KEY (`id`),
  KEY `k_1` (`k`)
) ENGINE=MyISAM AUTO_INCREMENT=100000001 DEFAULT CHARSET=latin1 |
```

CREATE STATUS

Name	Value
Name	sbtest1
Engine	MyISAM
Version	10
RowFormat	Fixed
Rows	100.00 m
AvgRowLength	189.00 Bytes
DataLength	17.60 GB
MaxDataLength	756.00 GB
IndexLength	1.70 GB

System Information

MySQL Summary

```
# Percona Toolkit MySQL Summary Report ######################
              System time | 2016-11-06 08:27:41 UTC (local TZ: CET +0100)
# Instances ################################################
  Port  Data Directory             Nice OOM Socket
  ===== ========================== ==== === ======
                                   0    0
# MySQL Executable #########################################
       Path to executable | /usr/sbin/mysqld
              Has symbols | No
# Report On Port 3306 ######################################
                     User | root@localhost
                     Time | 2016-11-06 09:27:41 (CEST)
                 Hostname | ps57
                  Version | 5.7.14-8-log Percona Server (GPL), Release 8, Revision 1f84ccd
                 Built On | Linux x86_64
                  Started | 2016-10-06 16:33 (up 30+16:54:37)
                Databases | 8
                  Datadir | /var/lib/mysql/
                Processes | 250 connected, 2 running
              Replication | Is not a slave, has 1 slaves connected
                  Pidfile | /var/run/mysqld/mysqld.pid (exists)
```

What happens on OS and Hardware Level

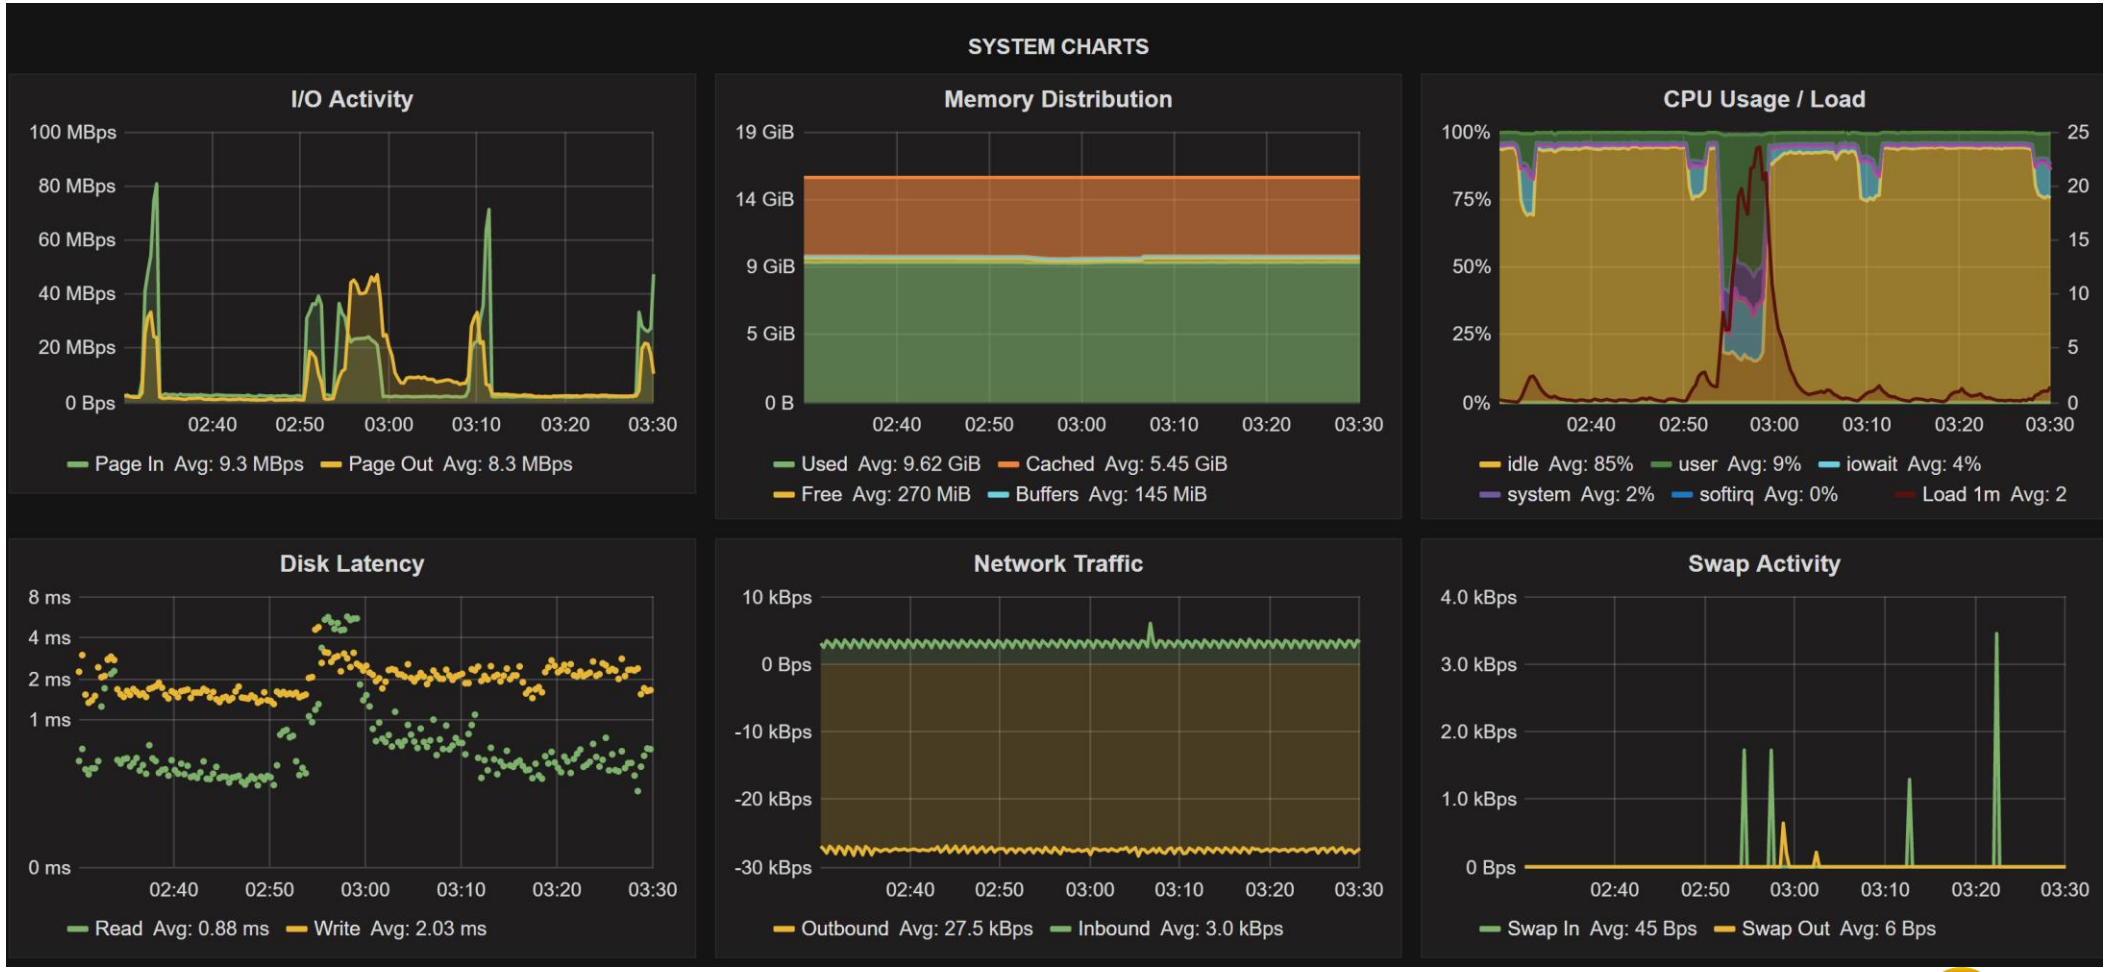

PERCONA

As well as Database Level

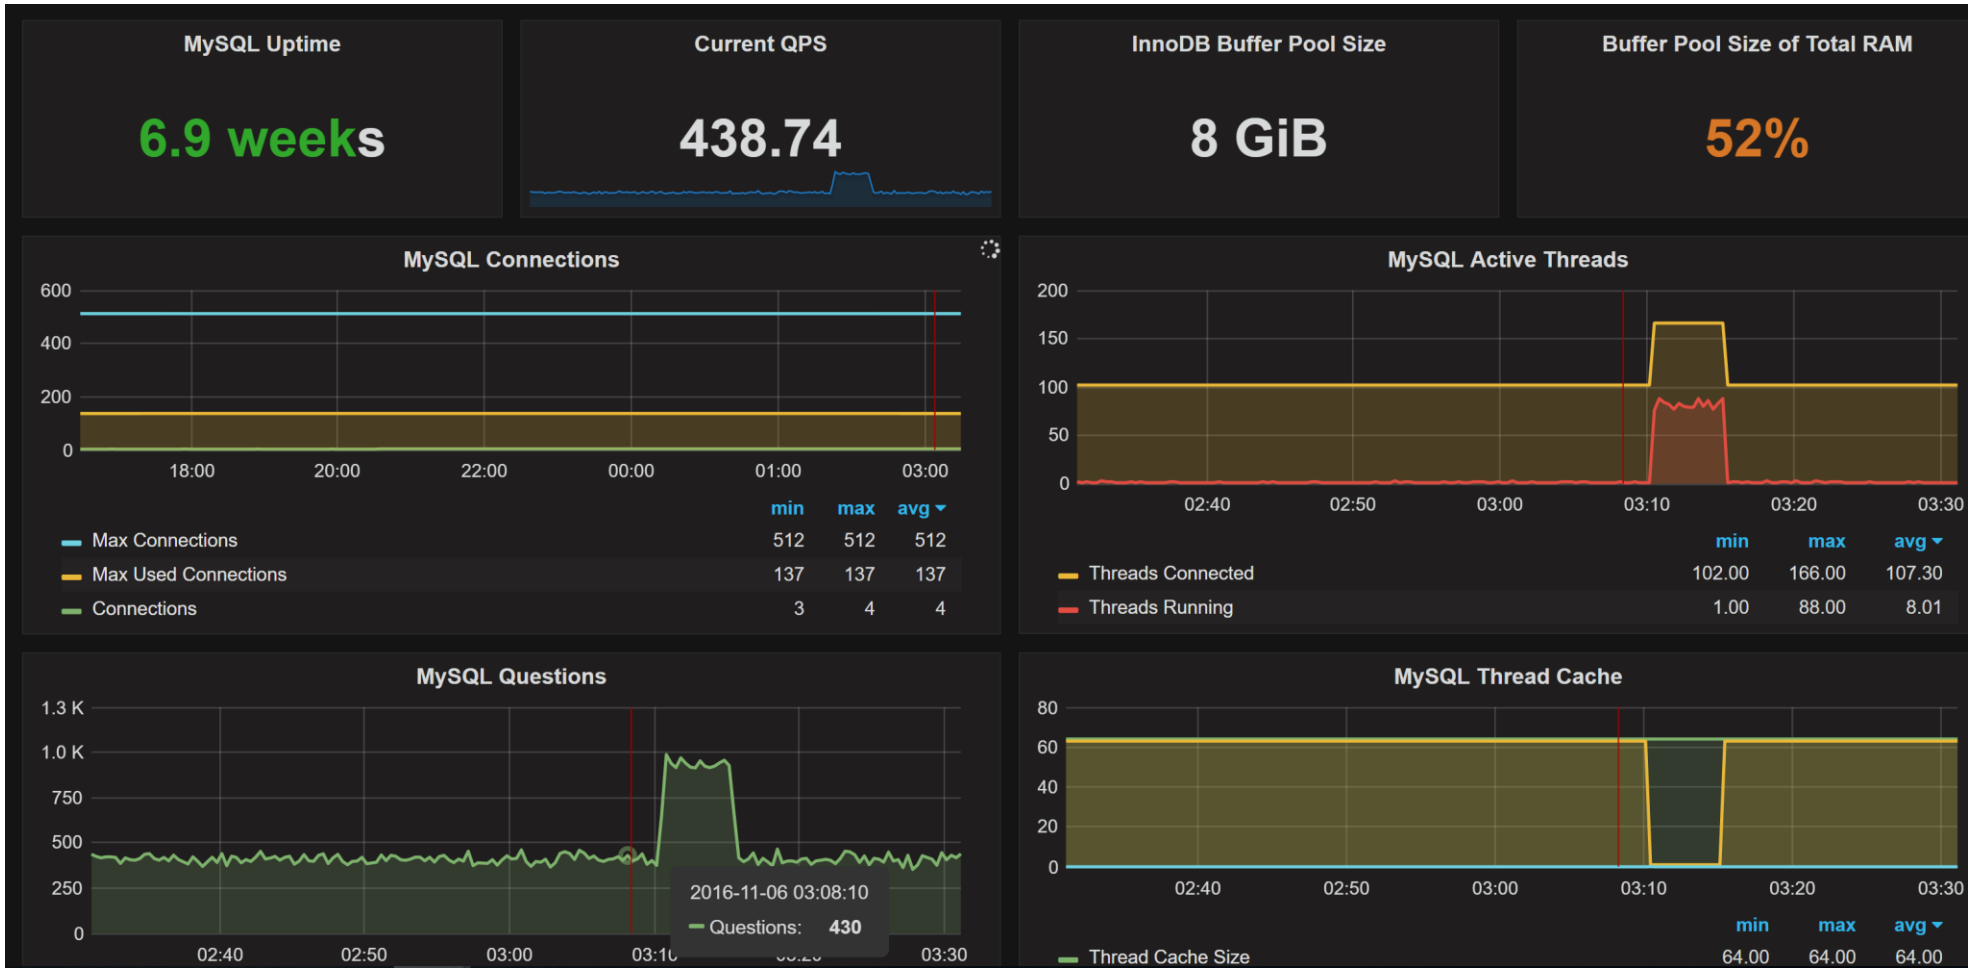

PERCONA

In-Depth MongoDB Dashboards

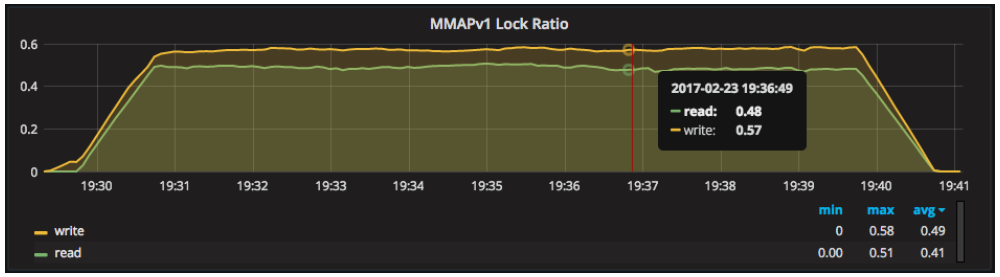

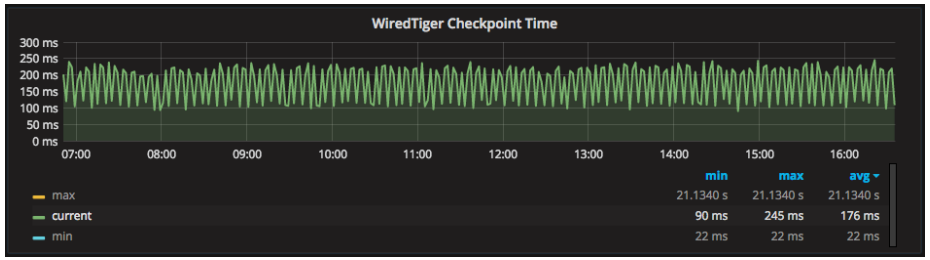

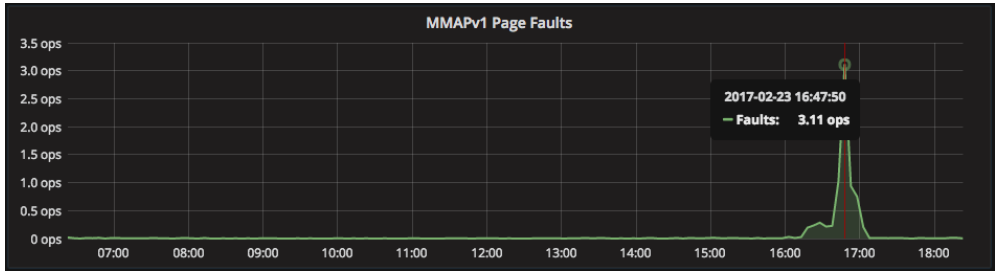

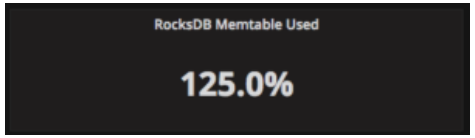

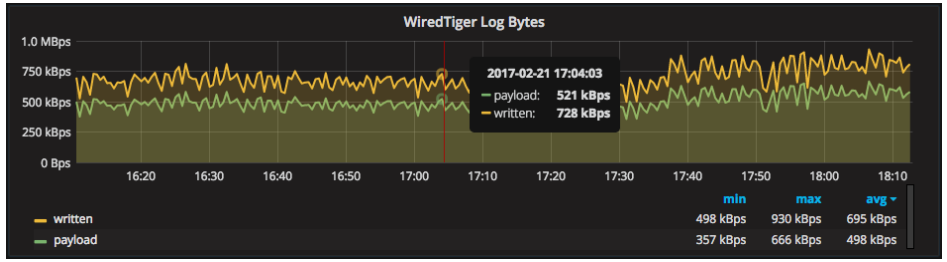

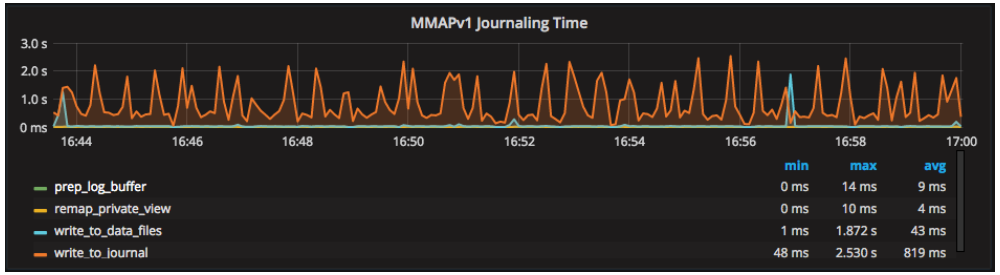

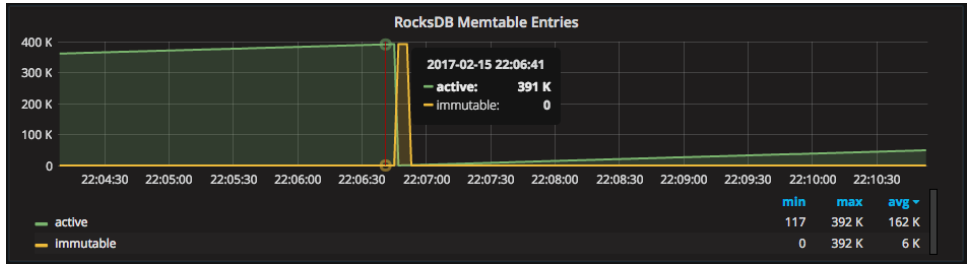

PERCONA

Best to check out Demo

http://pmmdemo.percona.com

ProxySQL

ProxySQL

- Traffic Management Solution for MySQL (MySQL Proxy)
- 100% Free ad Open Source
- Multiplexing
- Query Routing
- Sharding and Read Write Splitting
- Query Caching
- Integrates with Percona XtraDB Cluster and MySQL Group Replication
- http://www.proxysql.com/

Orchestrator

Orchestrator

- Visualize MySQL Replication
- Replication Monitoring
- Failover
- Change to Replication Topology
- Now Maintained by GitHub
- https://github.com/github/orchestrator

Orchestrator Screenshot

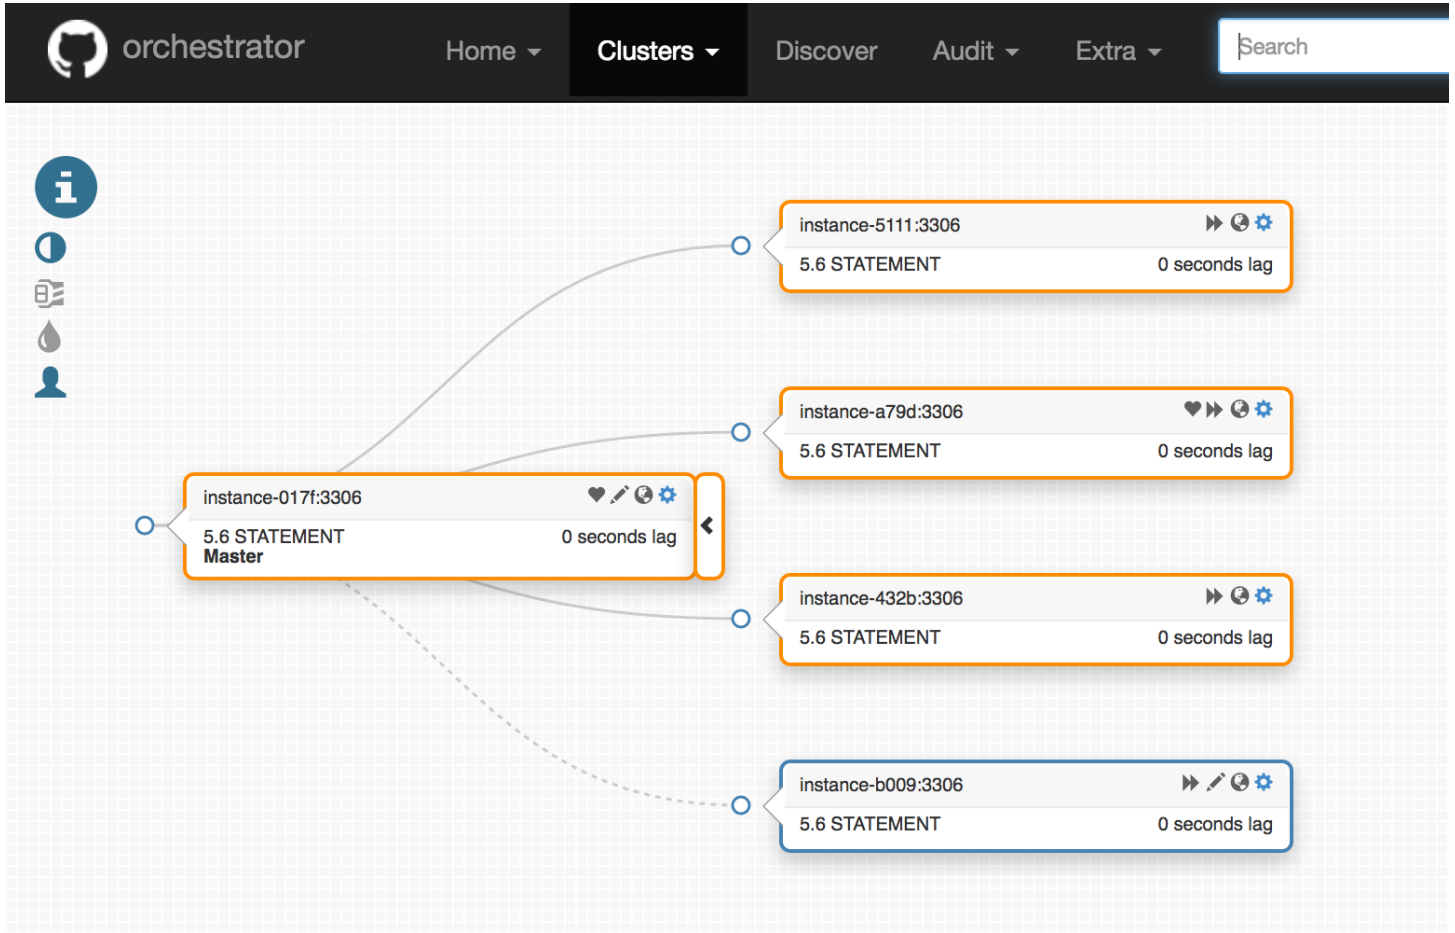

PERCONA

Gh-ost

GitHub's Online Schema Migrations for MySQL

Pt-online-schema-change on steroids

Does not use Triggers (less overhead and limitations)

Requires ROW binary log enabled

Tracks table changes through Binary Log

https://github.com/github/gh-ost

MongoDB 3.4

MongoDB 3.4 – New version most popular NoSQL Database

- Collation support for more than 100 languages and locales
- Decimal128
- $graphLookup operator for Graph Processing
- Faceted navigation
- Zones for Geographically Distributed Clusters
- Faster initial sync and better shard balancing
- Views

Percona Server for MongoDB 3.4

Percona Server for MongoDB 3.4

100% Compatible with MongoDB 3.4 Community Edition

Open Source with Alternatives to many MongoDB Enterprise Features

MongoRocks (RocksDB) and Percona Memory Engine

New: Sensitive Data Masking

New: Query Sampling

New: Hot Backup for WiredTiger and MongoRocks

Percona Memory Engine for MongoDB Benchmarks

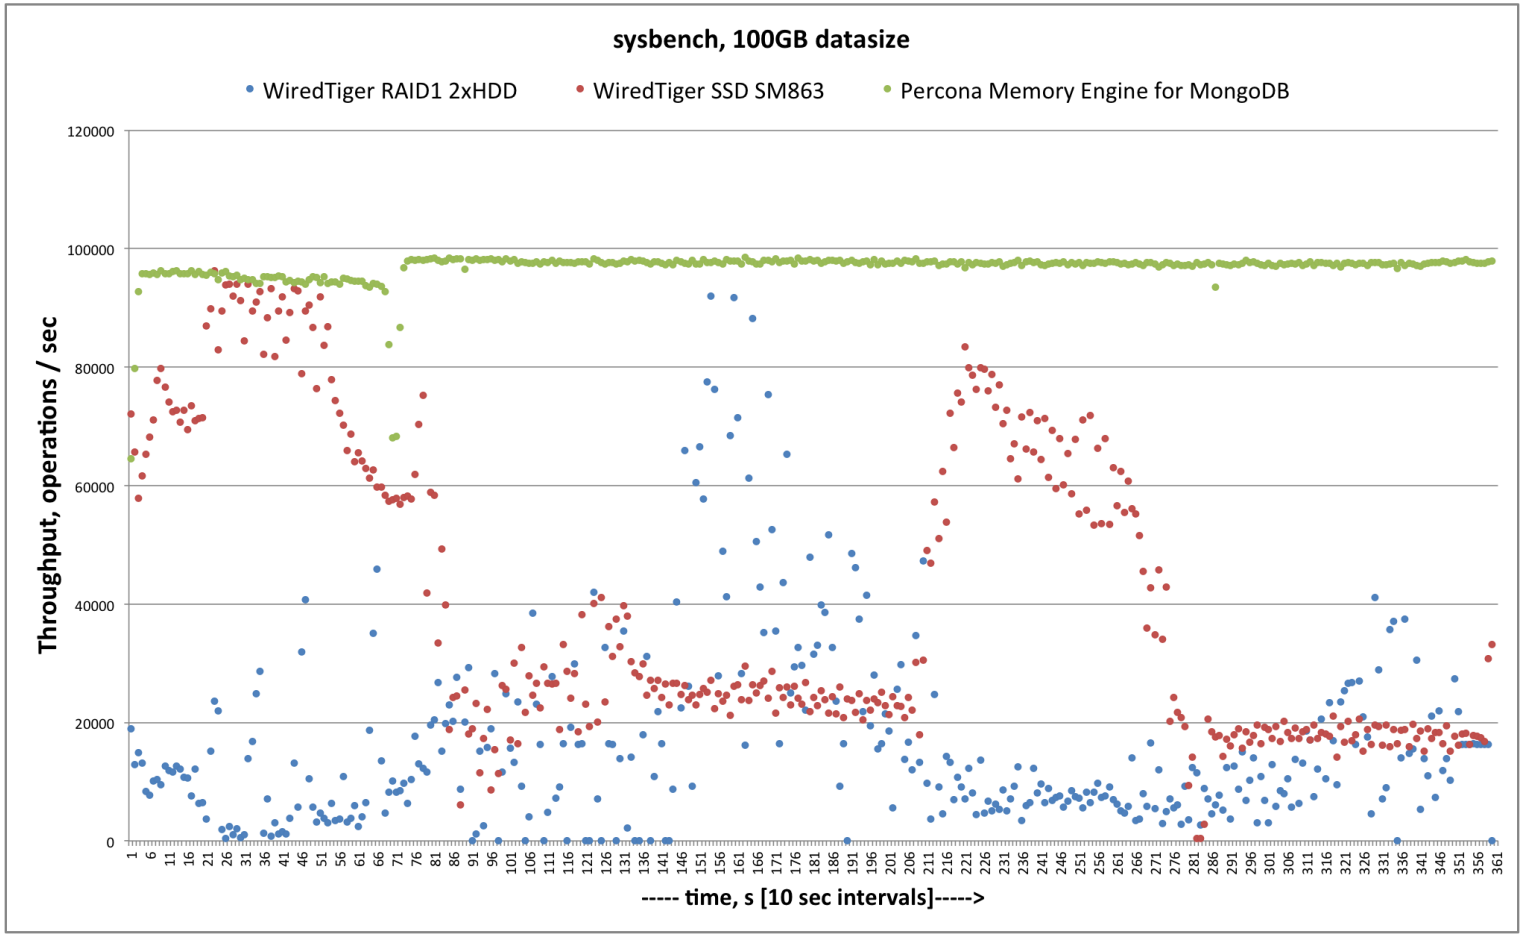

PERCONA

WiredTiger vs MongoRocks – write intensive

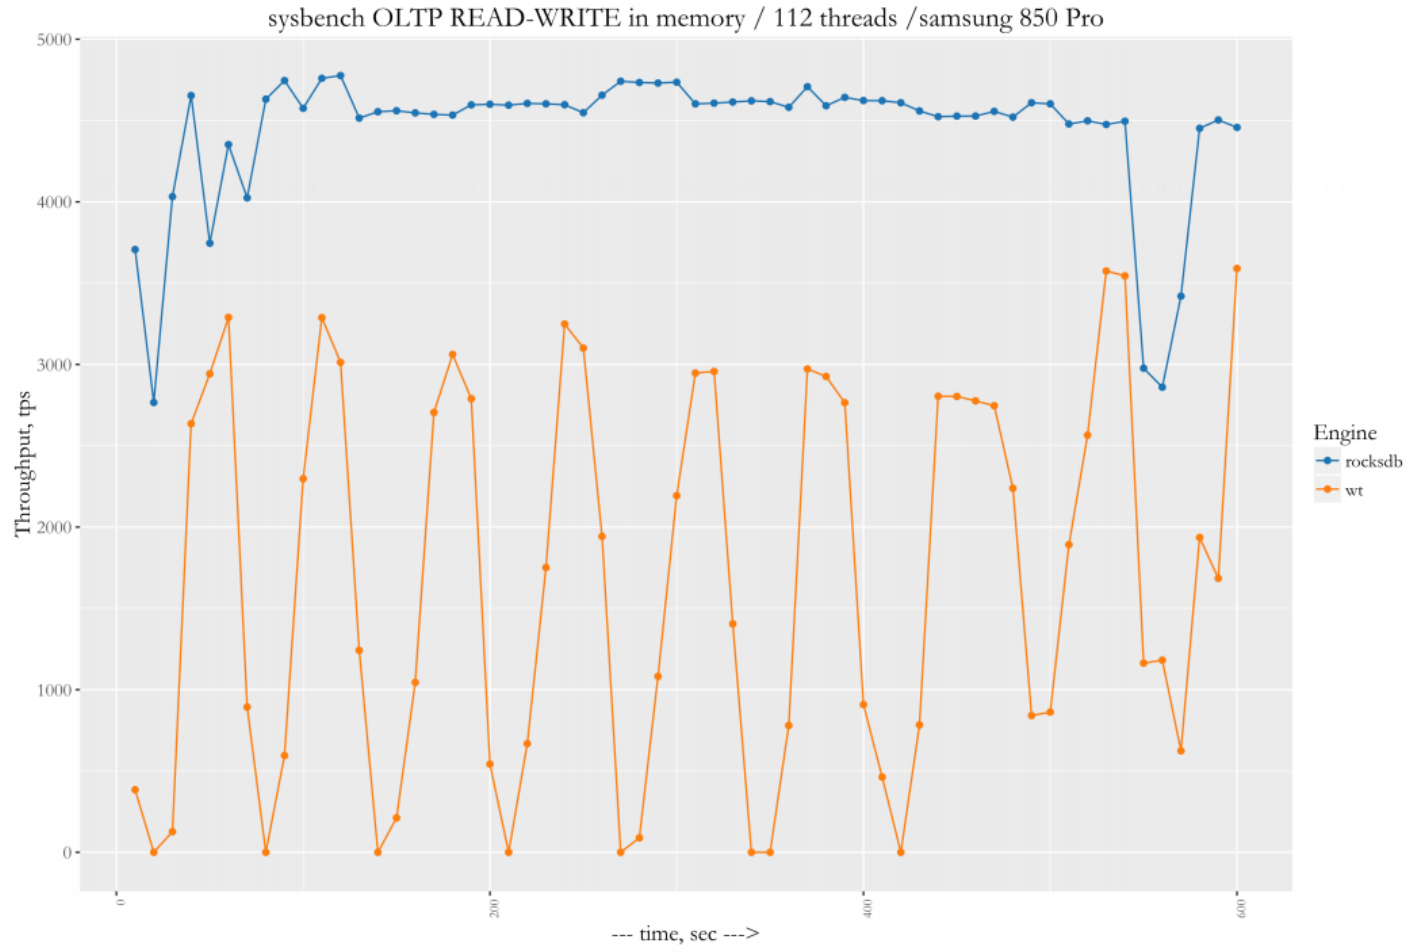

PERCONA

Before we take a Break…


BECOME A
PERCONA SUPERHERO
WE'RE HIRING
CONTACT
careers@percona.com
PERCONA

Have a Friend ?

Refer a friend and get $1000 if one is hired by Percona

… and eternal gratitude for helping to get the most amazing job

Percona Live Europe Call for Papers & Registration are Open!

Championing Open Source Databases

- MySQL, MongoDB, Open Source Databases
- Time Series Databases, PostgreSQL, RocksDB
- Developers, Business/Case Studies, Operations
- September 25-27th, 2017
- Radisson Blu Royal Hotel, Dublin, Ireland

Submit Your Proposal by July 17th!

www.percona.com/live/e17

Database Performance Matters