

# Troubleshooting with Percona Monitoring and Management

Overview, typical use cases and quick-start guide

---

**Dimitri Vanoverbeke & Michal Nosek**

Istanbul, October 4, 2019



# Agenda

---

1. Introduction
2. Performance monitoring tools, and PMM
3. MySQL, PostgreSQL, MongoDB - stories about problems
4. Solutions
5. PMM 2.0 Architecture
6. Quick-start guide
7. Q&A

# Survey

---


- Who uses more than one cloud provider?
- Who runs workloads both in the cloud and on-premise?
- Who runs more than one database?

# Our survey results:

---

- **41%** use more than one cloud provider
- **61%** run hybrid environments
- **92%** have more than one database in their environment

# Data growth pressure

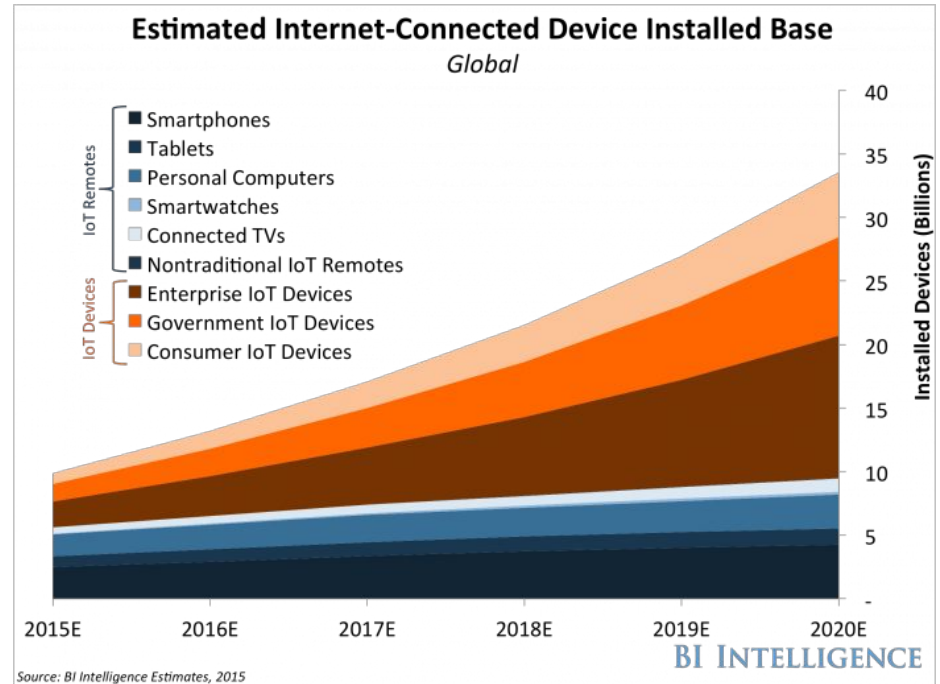


1 ZB =  $1000^7$  bytes =  $10^{21}$  bytes = 1000000000000000000000 bytes

Source: IDC DataAge 2025 whitepaper

# Performance pressure


- End-users expectations
- Complexity
- Requests volume
- Number of devices



# Agile development pressure

---

- Short release cycles
- Focus on fast delivery
- Different dev skills
- Unoptimized applications



Source: <https://www.techtrainees.com/relationship-between-devops-and-agile-development/>

# Tools to the rescue

---

## Database focused:

- VividCortex
- MySQL Enterprise Monitor
- MongoDB Ops Manager
- Amazon RDS Perf. Insights
- SeveralNines Cluster Control
- ...

## General Focused/Application

- DataDog
- Amazon CloudWatch
- NewRelic
- AppDynamics
- Zabbix
- ...



# Percona Monitoring and Management (PMM)

---

- Free and open-source
- MySQL, MongoDB, and PostgreSQL
- View historical data and trends
- Runs everywhere

# Problem stories



# Issue 1

---



## Industry: Logistics

A shipment database is being updated by drivers, keeping track of GPS locations, a driver tasks queue, completed tasks...

In the morning your drivers are receiving a tasks list, however instead of at 6:05, this list is taking longer and longer to be populated.

This list is populated by a batchjob triggered at 05:30



# Issue 1

## Impact:

Soon drivers will only be able to leave at 6:30 instead of 6:10.

They will be stuck in Peak traffic.

Putting the batch job earlier, provides the risk

That packages will only be shipped the next day



# Issue 1 - solution 1

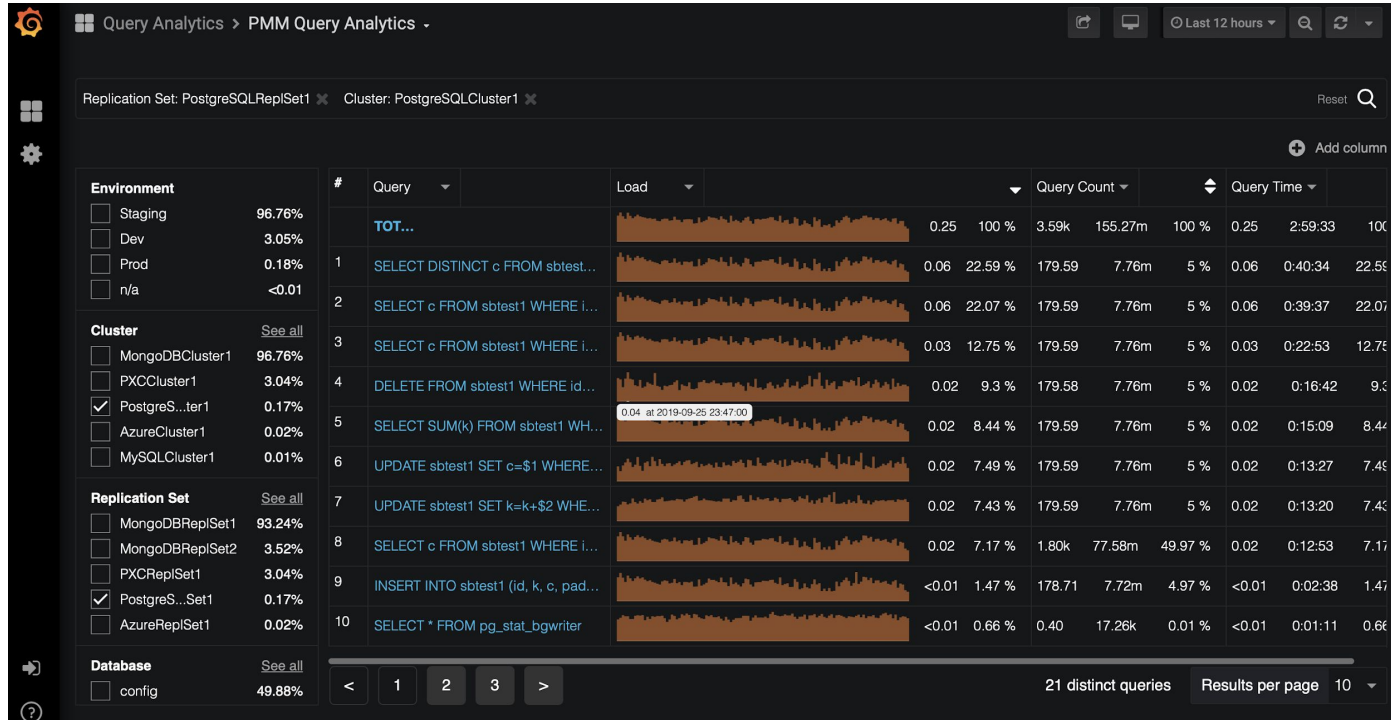
- Buy motorcycles and bikes to perform delivery.



- Invest in flying cars



# Issue 1 - solution



# Issue 1 - solution



```
DELETE FROM sbtest1 WHERE id=$1
```

49

Details Examples Explain Tables

Metric	Rate/Second	Total	Per Query Stats
Query Count <sup>?</sup>	100.33	4.33m 4.99% of total	1.00
Query Time <sup>?</sup>	0.02	0:16:21 12.98% of total	226.27 μs
Rows Sent <sup>?</sup>	99.75	4.31m 0.31% of total	0.99
Shared Blocks Dirtied <sup>?</sup>	0.37	15.85k 18.54% of total	0.00
Shared Block Cache Hits <sup>?</sup>	429.73	18.56m 1.15% of total	4.28

# Issue 1 - Solution

---

- **Query tuning is hard**
  - Without appropriate monitoring tools, reviewing query performance is less obvious.
- **Query monitoring is essential**
  - Optimal queries makes your database fly, what queries are taking the most resources or should be improved to free up resources
- **Query bottleneck identification**
  - Identifying bottlenecks can identify issues before they occur
- **Solving query issues, schema or system issues**



# Issue 2

---



**You are hosting a financial trading application, you are experiencing a sudden stop of transactions processing on the database. It's solving itself after a couple of seconds but during this period no financial transactions are happening.**

**This frustrates your clients requiring to be able to perform transactions reliably and at high speed**

# Issue 2

---



**Impact: You will lose business to a competitor capable of handling highload during trade peaks**

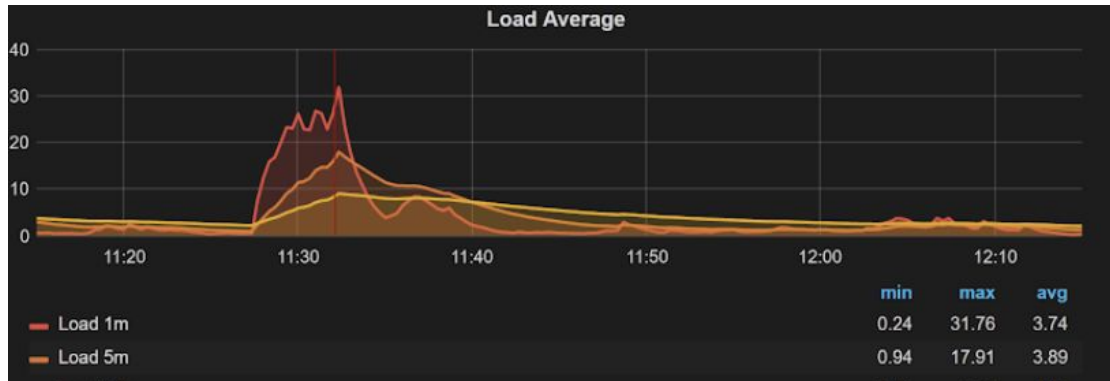



# Issue 2 - solution

- You go back to phone trading, because this high speed trading thing apparently is not workable in your organisation.



# Issue 2 - Identification



# Issue 2 - solution



# Issue 2 - Solution

---

- **Identification of a potential bottleneck is difficult without tools**
  - MySQL without monitoring/trending is like chicken without a waffle
  - Modification impact can be shown
- **Comparing server trending can show you tendencies**
  - Are you sending enough load to the replication hosts?
  - Is there a difference in between IO behaviour?
  - ...

# Issue 3

**Your company is offering document processing to the Police traffic department. Your application is storing millions of documents, however, when would be the moment to scale out your infrastructure and augment the MongoDB shards on your environment.**



# Issue 3



---

**Impact: Storing information, and accessing them for document creation will become slower and slower. Which means that offenders will get their fines later and later. You suddenly might detect a significant degradation in read/write performance.**

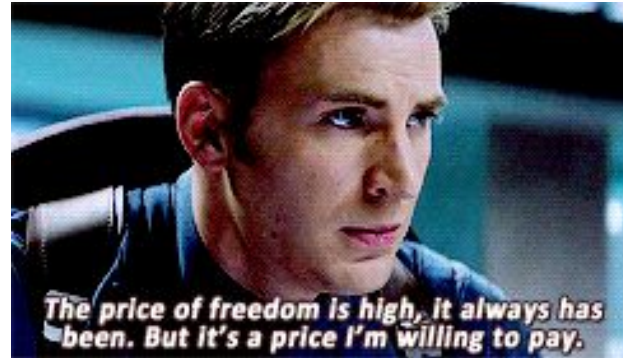
**The Justice department will be looking for guarantees or an SLA in processing...**




# Issue 3 - solution



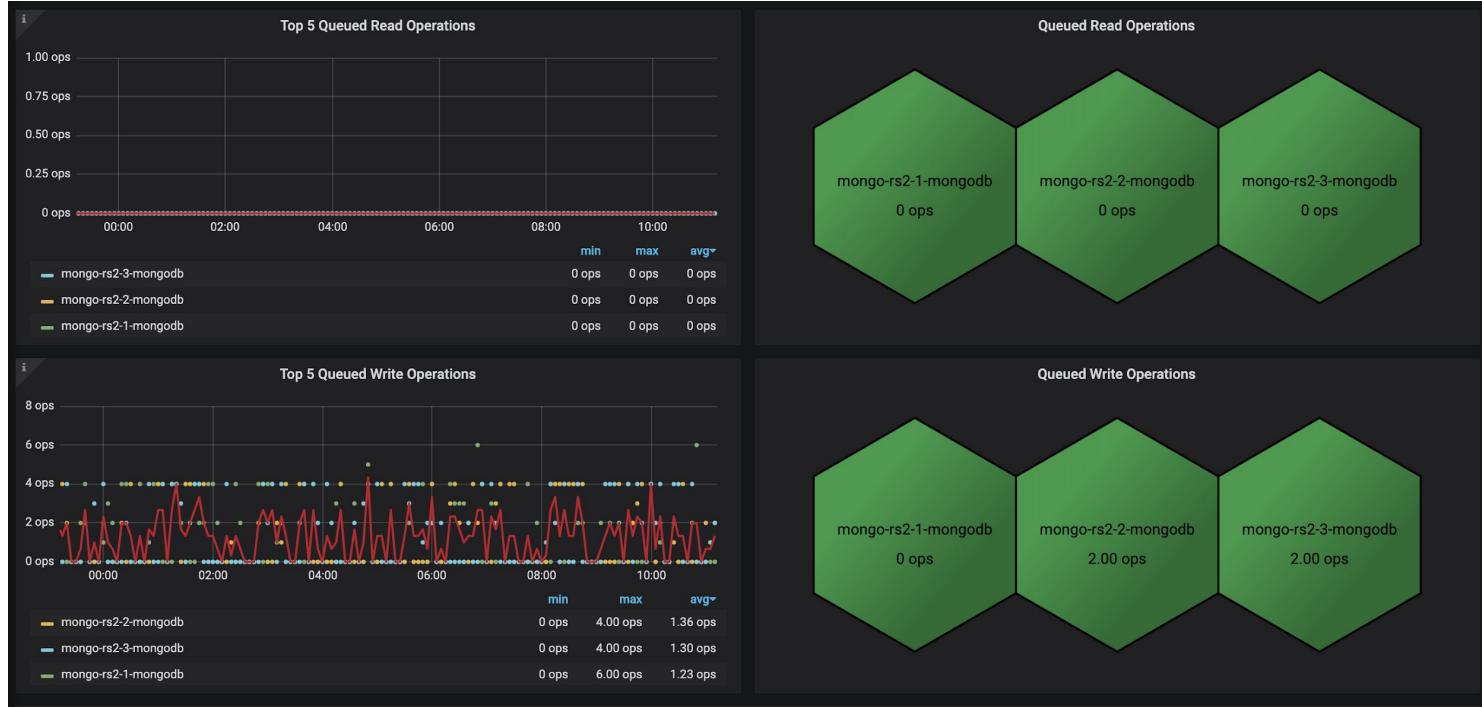
You advise the government to postpone traffic violation fines and make the world a free for all, we don't like to pay fines anyways.



# Issue 3 - solution



# Issue 3 - solution



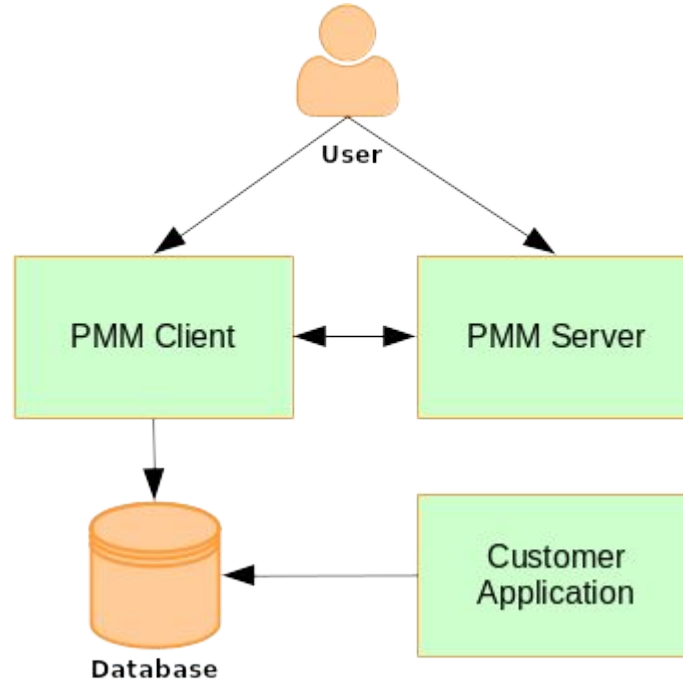
# Percona Monitoring and Management

---

Architecture and features

# Architecture - overview

---

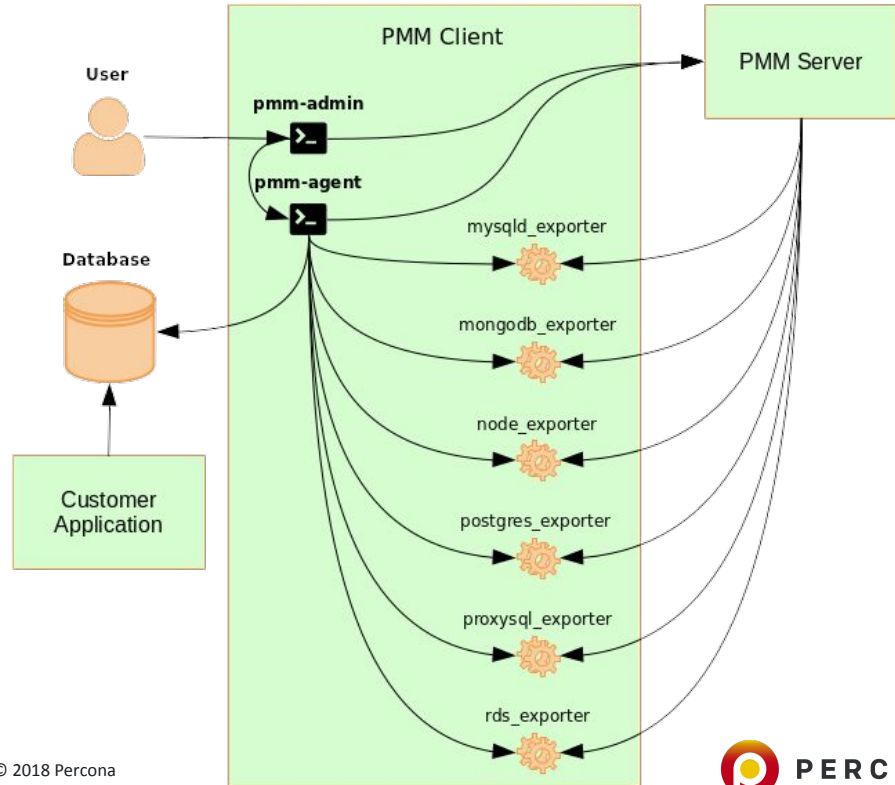


# Architecture - Client

- **Collects Server and database metrics, and query information**

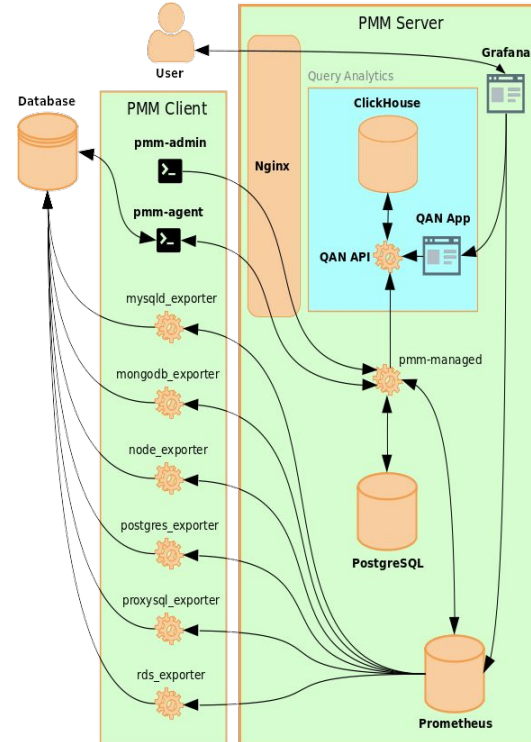
- **Components**

- Pmm-admin
- Pmm-agent
- exporters



# Architecture - Server

- **Query Analytics**
  - Grafana -> QAN APP
  - Using Clickhouse backend
- **Metrics Monitor**
  - Grafana <-> Prometheus
  - Prometheus aggregation



# Based on Open-source Technologies

---

- **Grafana (open source metric analytics & visualization suite)**
  - Visualizing time series data
- **Prometheus (Monitoring toolkit)**
  - Fetching data from clients (exporters)
- **Clickhouse (distributed column-oriented DBMS)**
  - Storing column based data



# Features summary

---



- **3 Dashboards**
- **Query analytics functionality**
- **PostgreSQL specific Dashboards**
  - Instance Summary
  - Instance Compare
  - Instance Overview



- **20 Dashboards**
- **Query analytics**
- **MYSQL specific Dashboards**
  - Amazon Aurora specific dashboard
  - MyRocks Details
  - Aria/MyISAM details
  - TokuDB details
  - PXC details
  - ...



- **9 Dashboards**
- **Query Analytics**
- **MongoDB specific dashboards**
  - instance details
  - Cluster overview
  - Wired Tiger details
  - Replset details
  - ...

# PMM2 Features summary

---

- Detailed **query analytics** and **filtering technologies** which enable you to identify issues faster than ever before.
- **Service-level dashboards** give you immediate access to the data you need.
- **PostgreSQL query tuning.**
- **Enhanced security protocols** to ensure your **data is safe.**
- Our new **API** allows you to **extend and interact** with third-party tools.

# Installation methods

---

- **PMM Server distributed 3 methods**
  - Docker (docker pull percona/pmm-server:2)
  - AMI (AWS Marketplace)
  - OVA (VMware, VirtualBox, etc)
- **pmm-client available as deb, RPM, linux binary, and source**
- **Marketplace for Amazon and Azure**
- **Easy documentation for setup on other cloud solutions**

# Quick-start guide

---

Demo time!

# Questions?

---

Documentation: <https://tinyurl.com/ya2kyf29>



**Champions of Unbiased  
Open Source Database Solutions**