



PERCONA
LIVE EUROPE
AMSTERDAM

30 Sept - 2 Oct 2019

CONFERENCE AND TUTORIAL **SCHEDULE**



Thanks to our sponsors:

Platinum



Gold



Silver



Community



This is the interactive guide to [Percona Live Europe 2019](#).
Links are clickable, including links back to the timetables at the bottom of the pages.

You can register at www.percona.com/live-registration

Sections

[Daily Schedules](#)

[Talks by Technology](#)

[Keynotes](#)

[Monday Tutorials](#)

[Tuesday Talks](#)

[Wednesday Talks](#)

[Speakers](#)

Spot talks by technology in the timetable: a color key

MySQL	MongoDB
PostgreSQL	Other Databases, Multiple Databases, and Other Topics

TUTORIALS DAY MONDAY, SEPTEMBER 30

	ROOM A	ROOM B	ROOM 7	ROOM 26 second floor	
9:00	<u>PostgreSQL For Oracle and MySQL DBAs and For Beginners</u>	<u>MySQL 101 Tutorial Part 1</u>	<u>MySQL 8.0 InnoDB Cluster: Easiest Tutorial!</u>	<u>Accelerating Application Development with Amazon Aurora</u>	9:00
12:00	LUNCH				12:00
1:30	<u>InnoDB Architecture and Performance Optimization Tutorial for MySQL 8</u>	<u>MySQL 101 Tutorial Part 2</u>	<u>Introduction to PL/pgSQL Development</u>	<u>Accelerating Application Development with Amazon Aurora</u>	1:30
4:30	WELCOME RECEPTION – EXPO HALL OPEN UNTIL 6PM				4:30
<hr/>					
	ROOM 8	ROOM 9	ROOM 10		
9:00	<u>MariaDB Server 10.4: The Complete Tutorial</u>	<u>Open Source Database Performance Optimization and Monitoring with PMM</u>	<u>Percona XtraDB Cluster Tutorial</u>		9:00
12:00	LUNCH				12:00
1:30	<u>Test Like a Boss: Deploy and Test Complex Topologies With a Single Command</u>	<u>A Journey with MongoDB HA. From Standalone to Kubernetes Operator</u>	<u>Getting Started with Kubernetes and Percona XtraDB Cluster</u>		1:30
4:30	WELCOME RECEPTION – EXPO HALL OPEN UNTIL 6PM				4:30

TUESDAY, OCTOBER 1

	ROOM A Building Large, Scalable & Secure Deployments	ROOM B Hot Topics, New Features & Trends	ROOM 7 Performance & Scalability	
9:00	<u>KEYNOTES</u>			9:00
10:30	COFFEE BREAK			10:30
11:00	<u>Our journey to better MySQL Availability Using Global Transaction IDs, ProxySQL and Consul</u>	<u>More Than a Query Language: SQL in the 21st Century</u>	<u>MySQL 8.0 Performance: Scalability & Benchmarks</u>	11:00
12:00	LUNCH			12:00
1:30	<u>Tarantool: A NoSQL Database with SQL</u>	<u>Percona XtraDB Cluster 8.0 (PXC-8.0)</u>	<u>What's New in PMM 2.0</u>	1:30
2:30	<u>How to protect PII data in MySQL While Allowing Customers to Access the Database</u>	<u>MyRocks and RocksDB Advanced Features and Performance</u>	<u>New Indexing and Aggregation Pipeline Capabilities in MongoDB 4.2</u>	2:30
3:30	COFFEE BREAK			3:30
4:00	<u>PCI/DSS Compliance with MySQL: 2019 Edition</u>	<u>What's New in Percona Server for MongoDB</u>	<u>Lessons From Building Automation for a Large Distributed Database</u>	4:00
5:00	<u>Backing up Wikipedia Databases</u>	<u>MySQL InnoDB Cluster: Advanced Configuration & Operation</u>	<u>InnoDB Management and Scalability Improvements in MySQL 8.0</u>	5:00
6:00	<u>From Scheduled Downtime to Self-Healing in Less Than a Year</u>	<u>MySQL 8.0.18: Latest Updates</u>	<u>Tracing and Profiling MySQL</u>	6:00
7:00	COMMUNITY EVENT BY BOOKING.COM AT HQ, HERENGRACHT 597, 1017CE AMSTERDAM			7:00

TUESDAY, OCTOBER 1

	ROOM 8 Public, Private & Hybrid Cloud	ROOM 9 Monitor, Manage, Maintain Databases @ Scale	ROOM 10 Reduce Costs & Complexity with OSS Databases	
9:00				9:00
10:30	COFFEE BREAK			10:30
11:00	<u>Vitess: Running Sharded MySQL on Kubernetes</u>	<u>Amazing Sandboxes with dbdeployer</u>	<u>Multi-Document Transactions Through MongoDB shell</u>	11:00
12:00	LUNCH			12:00
1:30	<u>MySQL on Google Cloud: The Good, The Bad, and The Ugly</u>	<u>MongoDB Data Security: Custom Rules and Views</u>	<u>Why PostgreSQL is Becoming A Migration Target in Large Enterprises</u>	1:30
2:30	<u>Comparison of Kubernetes Operators for MySQL</u>	<u>Pg_catalog Unrevealed! That Part of PostgreSQL You Are Probably Underusing</u>	<u>Enhancing MySQL Security</u>	2:30
3:30	COFFEE BREAK			3:30
4:00	<u>ClickHouse for Time-Series Real-time Analytics</u>	<u>Improving Enterprises HA and Disaster Recovery Solutions</u>	<u>Join Heterogeneous Databases Using PostgreSQL Foreign Data Wrappers</u>	4:00
5:00	<u>Deep Dive on Amazon Aurora</u>	<u>High Availability and Automatic Failover in PostgreSQL</u>	<u>VictoriaMetrics: Why and How We Built A Fast and Scalable Open Source Time Series Database</u>	5:00
6:00	<u>ClickHouse Features to Blow Your Mind</u>	<u>OpenCorporates: Providing Transparency for the Public Benefit</u>		6:00
7:00	COMMUNITY EVENT BY BOOKING.COM AT HQ, HERENGRACHT 597, 1017CE AMSTERDAM			7:00

WEDNESDAY, OCTOBER 2

	ROOM A Building Large, Scalable & Secure Deployments	ROOM B Hot Topics, New Features & Trends	ROOM 7 Performance & Scalability	
9:00	<u>Parted Ways with Partitioning? It's Time to Reconsider</u>	<u>Percona Server for MySQL 8.0: What It Is and How It Is Done</u>	<u>Billion Goods in Few Categories: How Histograms Save a Life?</u>	9:00
10:00	<u>Percona Backup for MongoDB</u>	<u>Databases, Crypto & Decentralization</u>	<u>ClickHouse Materialized Views: A Secret Weapon for High Performance Analytics</u>	10:00
10:30	COFFEE BREAK			10:30
11:00	<u>Data Protection and OSS in the Age of GDPR</u>	<u>What's New on Sharding in MongoDB 4.2</u>	<u>My First 90 Days With Vitess</u>	11:00
11:30	<u>Securing MySQL with Vault and Table Level Encryption</u>	<u>Benchmarking Should Never be Optional</u>	<u>10 Common Mistakes (Java) Developers Make When Writing SQL</u>	11:30
12:00	LUNCH			12:00
1:30	<u>Strength in Numbers: Slack's Database Architecture</u>	<u>MySQL 8.0: The New Replication Features</u>	<u>Optimize and Troubleshoot MySQL Using PMM 2.0</u>	1:30
2:30	<u>Managing MySQL at Scale in Facebook</u>	<u>Graph Databases: Introduction, Standardization, Opportunities</u>	<u>How to Instrument Your Code in Performance Schema</u>	2:30
3:30			<u>gdb Basics for MySQL DBAs</u>	3:30
4:00	<u>KEYNOTES, PRIZE GIVING, AND SUMMING UP</u>			4:00

WEDNESDAY, OCTOBER 2

	ROOM 8 Public, Private & Hybrid Cloud	ROOM 9 Monitor, Manage, Maintain Databases @ Scale	ROOM 10 Reduce Costs & Complexity with OSS Databases	
9:00	<u>MySQL Has Gone Away: An In-Depth Look at the MySQL Networking Implementation</u>	<u>How a Modern Database Gets Your Data Fast: MariaDB Query Optimizer</u>	<u>Maintenance for MongoDB Replica Sets</u>	9:00
10:00	<u>Fortify Your MySQL Data Security in AWS Using ProxySQL Firewall</u>	<u>Percona Distribution for PostgreSQL</u>	<u>MySQL Shell : The Best DBA tool? Extend the Shell with the New Extension Infrastructure</u>	10:00
10:30	COFFEE BREAK			10:30
11:00	TBC	<u>Running PMM in Production at Tessi</u>	<u>Handling Transaction ID Wraparound in PostgreSQL</u>	11:00
11:30	<u>Top 10 Mistakes When Migrating From Oracle to PostgreSQL</u>	<u>Large Scale Deployment of SSL/TLS For MySQL</u>	<u>Protecting your Secrets</u>	11:30
12:00	LUNCH			12:00
1:30	<u>MongoDB Analysis with Prometheus and Grafana</u>	<u>How to Upgrade Like a Boss to MySQL 8.0?</u>	<u>Running Elasticsearch at Large Scale</u>	1:30
2:30	<u>Becoming Cloud Native: How Percona Brings Databases to Kubernetes Using Operators</u>	<u>The Database is Broken. Now What?</u>	<u>Automatic Upgrade and New Error Logging in MySQL 8.0</u>	2:30
3:30	<u>PostgreSQL Plan at Execution Time: A Quick Show</u>	<u>JSON Array Indexes in MySQL</u>		3:30
4:00				4:00

MySQL & MariaDB

MONDAY

[MySQL 101 Tutorial Parts 1 & 2](#)

[MySQL 8.0 InnoDB Cluster: Easiest Tutorial!](#)

[MariaDB Server 10.4: The Complete Tutorial](#)

[Percona XtraDB Cluster Tutorial](#)

[InnoDB Architecture and Performance Optimization Tutorial for MySQL 8](#)

[Getting Started with Kubernetes and Percona XtraDB Cluster](#)

[Test Like a Boss: Deploy and Test Complex Topologies With a Single Command](#)

TUESDAY

[Our journey to better MySQL Availability Using Global Transaction IDs, ProxySQL and Consul](#)

[MySQL 8.0 Performance: Scalability & Benchmarks](#)

[Vitess: Running Sharded MySQL on Kubernetes](#)

[Amazing Sandboxes with dbdeployer Percona XtraDB Cluster 8.0 \(PXC-8.0\)](#)

[MySQL on Google Cloud: The Good, The Bad, and The Ugly](#)

MySQL & MariaDB

TUESDAY

[How to protect PII data in MySQL While Allowing Customers to Access the Database](#)

[Comparison of Kubernetes Operators for MySQL](#)

[Enhancing MySQL Security](#)

[PCI/DSS Compliance with MySQL: 2019 Edition](#)

[Lessons From Building Automation for a Large Distributed Database](#)

[Improving Enterprise HA and Disaster Recovery Solutions](#)

[Backing up Wikipedia Databases](#)

[MySQL InnoDB Cluster: Advanced Configuration & Operation](#)

[InnoDB Management and Scalability Improvements in MySQL 8.0](#)

[From Scheduled Downtime to Self-Healing in Less Than a Year](#)

[MySQL 8.0.18: Latest Updates](#)

[Tracing and Profiling MySQL](#)

[OpenCorporates: Providing Transparency for the Public Benefit](#)

MySQL & MariaDB

WEDNESDAY

[Percona Server for MySQL 8.0: What It Is and How It Is Done](#)

[Billion Goods in Few Categories: How Histograms Save a Life?](#)

[MySQL Has Gone Away: An In-Depth Look at the MySQL Networking Implementation](#)

[How a Modern Database Gets Your Data Fast: MariaDB Query Optimizer](#)

[My First 90 Days With Vitess](#)

[Fortify Your MySQL Data Security in AWS Using ProxySQL Firewall](#)

[MySQL Shell : The Best DBA tool? Extend the Shell with the New Extension Infrastructure](#)

[Securing MySQL with Vault and Table Level Encryption](#)

[Large Scale Deployment of SSL/TLS For MySQL](#)

[MySQL 8.0: The New Replication Features](#)

TECHNOLOGIES

MySQL & MariaDB

WEDNESDAY

[Optimize and Troubleshoot MySQL Using PMM 2.0](#)

[How to Upgrade Like a Boss to MySQL 8.0?](#)

[Managing MySQL at Scale in Facebook](#)

[How to Instrument Your Code in Performance Schema](#)

[The Database is Broken. Now What?](#)

[Automatic Upgrade and New Error Logging in MySQL 8.0](#)

[gdb Basics for MySQL DBAs](#)

[JSON Array Indexes in MySQL](#)

MongoDB

MONDAY

[A Journey with MongoDB HA. From Standalone to Kubernetes Operator](#)

TUESDAY

[Multi-Document Transactions Through MongoDB shell](#)

[MongoDB Data Security: Custom Rules and Views](#)

[New Indexing and Aggregation Pipeline Capabilities in MongoDB 4.2](#)

[What's New in Percona Server for MongoDB](#)

WEDNESDAY

[Maintenance for MongoDB Replica Sets](#)

[Percona Backup for MongoDB](#)

[What's New on Sharding in MongoDB 4.2](#)

[MongoDB Analysis with Prometheus and Grafana](#)

PostgreSQL

MONDAY

[PostgreSQL For Oracle and MySQL DBAs and For Beginners](#)

[Introduction to PL/pgSQL Development](#)

TUESDAY

[Why PostgreSQL is Becoming A Migration Target in Large Enterprises](#)

[Pg_catalog Unrevealed! That Part of PostgreSQL You Are Probably Underusing](#)

[Join Heterogeneous Databases Using PostgreSQL Foreign Data Wrappers](#)

[High Availability and Automatic Failover in PostgreSQL](#)

[SEE WEBSITE FOR LATE CHANGES](#)

PostgreSQL

WEDNESDAY

[Parted Ways with Partitioning? It's Time to Reconsider](#)

[Percona Distribution for PostgreSQL](#)

[Handling Transaction ID Wraparound in PostgreSQL](#)

[Top 10 Mistakes When Migrating From Oracle to PostgreSQL](#)

[PostgreSQL Plan at Execution Time: A Quick Show](#)

Other Databases, Multiple Databases & Other Technologies

MONDAY

[Open Source Database Performance Optimization and Monitoring with PMM Accelerating Application Development with Amazon Aurora](#)

TUESDAY

[More Than a Query Language: SQL in the 21st Century](#)

[Tarantool: A NoSQL Database with SQL](#)

[What's New in PMM 2.0](#)

[MyRocks and RocksDB Advanced Features and Performance](#)

[ClickHouse for Time-Series Real-Time Analytics](#)

[Deep Dive on Amazon Aurora](#)

[VictoriaMetrics: Why and How We Built A Fast and Scalable Open Source Time Series Database](#)

[ClickHouse Features to Blow Your Mind](#)

Other Databases, Multiple Databases & Other Technologies

WEDNESDAY

[Databases, Crypto & Decentralization](#)

[ClickHouse Materialized Views: A Secret Weapon for High Performance Analytics](#)

[Data Protection and OSS in the Age of GDPR](#)

[Running PMM in Production at Tessi](#)

[Benchmarking Should Never be Optional](#)

[10 Common Mistakes \(Java\)Developers Make When Writing SQL](#)

[Protecting your Secrets](#)

[Strength in Numbers: Slack's Database Architecture](#)

[Graph Databases: Introduction, Standardization, Opportunities](#)

[Running Elasticsearch at Large Scale](#)

[Graph Databases: Introduction, Standardization, Opportunities](#)

[Becoming Cloud Native: How Percona Brings Databases to Kubernetes Using Operators](#)

KEYNOTE ADDRESSES TUESDAY, OCTOBER 1

9:00AM	Matt Yonkovit, CXO, Percona	Welcome to Percona Live Europe, 2019
9:05AM	Peter Zaitsev CEO, Percona	The State of Open Source Databases Peter explores the current topics and trends impacting the database ecosystem, as well as the businesses that run open-source databases as part of their infrastructure. Will cloud-native surpass public cloud usage? How is the changing licensing around open source going to impact you? What keeps most of your peers up at night? What will impact your business in 2020 and beyond?
9:30AM	Geir Høydalsvik, Software Development Director, Oracle	The State of the Dolphin Where are we with MySQL 8.0 one year after GA? How was the journey and how the Community impacted the road map? During this session come to discover our new release way, allowing us to be even closer to users looking for performance, stability and innovation.
9:55AM	John Boero, Channel Solutions Engineer, Hashicorp	Unlocking the Cloud Operating Model To unlock the fastest path to the value of the cloud, enterprises must consider how to industrialize the application delivery process across each layer of the cloud: embracing the cloud operating model, and tuning people, processes, and tools to it. The essential implication of the transition to the cloud is the shift from static infrastructure to dynamic infrastructure: from a focus on configuration, and management of a static fleet of IT resources, to a focus on provisioning, securing, connecting, and running dynamic resources on demand. In this session, we'll discuss what the Cloud Operating Model means for your enterprise.
10:25AM	Matt Yonkovit, CXO, Percona	Announcements and Release to Sessions

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

KEYNOTE ADDRESSES WEDNESDAY, OCTOBER 2

4:00PM	Matt Yonkovit, CXO, Percona	Welcome Back to Percona Live Europe, 2019
4:05PM	Matt Yonkovit, CXO, Percona	Why Have Efforts to Simplify Led to More Complexity? <i>How to Survive What has Become an Increasingly Complex and Confusing World.</i> <p>A great deal of time and effort goes into making things easier, from simplification through automation, DevOps processes, and the emergence of new technology. Despite the best efforts of our brightest minds, systems still crash, databases are breached (seemingly daily), and smart people still do stupid things. Despite our best efforts, advancements have led to more complexity, less unity, and more confusing options than ever before. Why did this happen and how can we address it?</p>
4:20PM	Craig Kerstiens, Head of product for Citus Cloud, Microsoft	Postgres: From Respected by DBAs to Loved by Application Developers <p>Ten years ago, Postgres was a safe and stable database, although not at the same level of user-friendliness or rich feature set as other databases. But, that all started to change as it learned from MySQL and Mongo. We'll examine some of that journey: how it went from being respected by DBAs to beloved by application developers. We'll close with a look at what the next five years may hold.</p>
4:45PM	Panel Discussion	Business Challenges Facing Modern Database Deployments <p>We will engage with business leaders from around the globe and ask them about trends, changes, and issues they face day-to-day. We will ask them for their advice on how to build scalable, secure, and world-class applications that meet user needs.</p>
5:10PM	Matt Yonkovit, CXO, Percona	Wrap Up: See You At Percona Live 2020!

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

TUTORIAL DAY
MONDAY, SEPTEMBER 30

PostgreSQL For Oracle and MySQL DBAs and For Beginners

Avinash Vallarapu, Percona

This tutorial is focused on teaching PostgreSQL to professionals who are already experienced with MySQL or Oracle or SQL Server. Relational databases share a lot of similar concepts. You may find it easier when the concepts of PostgreSQL are related with Oracle/MySQL. Thus, I would start with PostgreSQL Architecture for Oracle/MySQL DBAs.

As we discuss from basics, even a beginner could adapt the concepts easily. It starts with an overview of the PostgreSQL architecture and a hands-on lab, helping the audience apply the concepts presented, starting with a simple and quick installation of PostgreSQL and going all the way to setup and maintain a standby (replica).

By the end of this tutorial, attendees will have a good understanding of how PostgreSQL works and feel confident enough about it.

Following topics are covered in this tutorial.

1. Brief history of PostgreSQL and its features
2. PostgreSQL Architecture for MySQL/Oracle DBAs and beginners.
3. Installation of PostgreSQL
4. Creating/Start/Stop a PostgreSQL Cluster and the components of a Data Directory
5. PostgreSQL parameters/settings you need to know
6. PostgreSQL User Management
7. MVCC in PostgreSQL
8. VACUUM and Transaction ID Wraparound in PostgreSQL
9. Types of Indexes
10. EXPLAIN in PostgreSQL
11. Logical and Physical Backups
12. Types of Replication in PostgreSQL
13. Streaming replication and PITR (Point-in-time-recovery)

MySQL 101 Tutorial Parts 1 & 2

Tom de Cooman, Dim0 Vanoverbeke, Michal Nosek, Percona

This full-day tutorial will be mainly hands-on. We'll show some slides on each of the various topics presented, but the goal is to have YOU the student work in a MySQL environment as much as possible.

What will we cover? Almost every aspect of being a MySQL DBA, on MySQL 5.7:

- We will install MySQL and configure the right privileges for your system
- Once we have a database running, we will get a better understanding of what our database is doing and look at the different possible diagnostics and monitoring options in MySQL using PMM and other tools.
- We will troubleshoot issues
- We will discuss and try some of the different backup options, since not losing any data is an important goal for a MySQL DBA
- We will clarify and troubleshoot replication

MySQL 8.0 InnoDB Cluster: Easiest Tutorial!

Frédéric Descamps, Oracle

During this tutorial we will show how easy it's now to create, manage and monitor a MySQL InnoDB Cluster. HA for MySQL has never been so easy. We will migrate an existing environment with minimal downtime and show different tuning settings related to Group Replication. The audience will also get familiar with the MySQL Shell.

MariaDB Server 10.4: The Complete Tutorial

Colin Charles, Consultant

Let's take a simple Oracle-backed application, and migrate it to MariaDB Server 10.4. While we are at it, let's learn about how to use system versioned tables, Galera 4, the MyRocks storage engine, Flashback to rollback tables/databases to an older snapshot, and also use the built-in threadpool.

This and more is possible in 3 hours, as MariaDB Server 10.4 has been out for sometime now (June 2019) and it has many new features, some of which MySQL does not have. Feature wise, it is important to know what MariaDB Server 10.4 has (e.g. system tables in the Aria storage engine, ability to reload SSL certificates without a restart and more!) and what it lacks compared to MySQL 8.0 (group replication, the X Protocol, etc.).

All participants (developer or operations personnel) will leave more knowledgeable on how to better manage, observe, and secure their MariaDB Servers. In addition, there will be a focus on High Availability as well as backups (covering Mariabackup) & disaster recovery.

Attendees will field fast-paced practice sessions on:

- Upgrading their MySQL servers to MariaDB successfully.
- Learning the features that MariaDB offers beyond MySQL.
- Understanding the compatibility between MariaDB and MySQL, including how applications will work and react to it.
- Have an overview of storage engines, but focus on what the most popular ones can be used for.
- Understanding MariaDB Server replication options

The structure of this tutorial will be practical, so please bring a laptop, ensure you have the ability to run either SSH to a remote machine, or the ability to run "dbdeployer" on Linux. There will be virtual machines provided as an alternative.

Open Source Database Performance Optimization and Monitoring with PMM

Michael Coburn, Vinicius Grippa, Sveta Smirnova, Carlos Tutté, Percona

This hands-on tutorial covers how to set up monitoring for MySQL, MongoDB, and PostgreSQL database servers using the Percona Monitoring and Management (PMM) platform.

PMM is an open-source platform for managing and monitoring MySQL, MongoDB, and PostgreSQL performance. It provides thorough time-based analysis for database servers to ensure that they work as efficiently as possible. You will learn:

- MySQL, MongoDB, and PostgreSQL monitoring best practices
- Metrics and time series
- Data collection, management and visualization tools
- Monitoring deployment
- How to use graphs to spot performance issues
- Query analytics for MySQL and MongoDB, and now PostgreSQL!
- Alerting using Grafana
- Trending and capacity planning
- How to monitor HA

Please bring a laptop with an ssh client as you will be connecting to a cloud-provisioned instance on which you will deploy PMM 2.

Percona XtraDB Cluster Tutorial

Tibor Köröcz, Percona

Percona XtraDB Cluster (PXC) is the leading HA solution for MySQL with its write-anywhere, synchronous replication capabilities. This half-day tutorial will introduce you to PXC, explain core concepts, discuss advanced techniques, cover common “gotchas”, and utilize ProxySQL to create the most highly-available MySQL cluster solution currently available.

InnoDB Architecture and Performance Optimization Tutorial for MySQL 8

Peter Zaitsev, Percona

MySQL 8 is the most commonly used storage engine for MySQL and Percona Server. It is the focus for the majority of storage engine development and Percona Server teams.

In this tutorial, we will look at the MySQL 8 architecture, including new developments in MySQL 8 as well as Percona Server. We will explain how it should be used to get the best performance from your application.

We will provide specific advice on server configuration, schema design, application architecture, and hardware choices.

This tutorial is updated from its previous version, to cover new MySQL 8 and Percona Server features in more detail.

Introduction to PL/pgSQL Development

Jim Mlodgenski AWS

When developing against a PostgreSQL database, many times you will use SQL from your applications, but sometimes you need more. You might need to run some of your logic on the database server for performance or you may want to just create some server-side functions to simplify your application development.

There is an infinite number of reasons why you would want to write some code that runs on the PostgreSQL server and PostgreSQL has a rich set of options when you want to do that, but most times, you will use PL/pgSQL. The primary objective of this tutorial is to introduce you to the PostgreSQL PL/pgSQL programming language and to give you the practical skills to be a productive developer working with PostgreSQL databases.

In the course of this tutorial, you will be introduced to a variety of fundamental PL/pgSQL functions like:

- How to declare and use PL/SQL Variables
- How to create functions and procedures
- How to write triggers
- Handle run time errors

And much more...

Test Like a Boss: Deploy and Test Complex Topologies With a Single Command

Giuseppe Maxia, dbdeployer

Using MySQL databases or its derivatives nowadays is much more than running a main server in production.

It's much more than maintaining a master/slave or group replication system.

New versions and new features are coming at great speed, and users need to try the novelties quickly, possibly in their in-house hardware. dbdeployer allows users to deploy simple and complex topologies with a single command, using a single host, without interfering with existing servers in the same machine.

In this tutorial you will learn how to install a test cluster with asynchronous replication, with group replication (single primary and multi primary), with NDB cluster, or other multi-master topologies such as fan-in and all-masters.

We will also see how to replicate from single server to group replication, from group replication to master/slave, or from two NDB clusters.

Also we'll see how to download and deploy in a few commands, without need to use a browser, but just with dbdeployer's assistance.

Finally, we'll see how to import existing deployments into a sandbox and how to clone a sandbox into another.

Getting Started with Kubernetes and XtraDB Cluster

Matthew Boehm, Percona

Have you ever wanted to properly learn how to run Percona XtraDB Cluster on top of Kubernetes?

Join us in this 3 hour tutorial where we cover the quick basics of containers, and then dive right in to the Percona Kubernetes Operator for XtraDB Cluster.

You'll learn:

- How to install the operator, create a highly available production-ready XtraDB Cluster.
- How to manage backups and recovery.
- How to scale your cluster up and down to meet demand.

A Journey with MongoDB HA From Standalone to Kubernetes Operator – A Hands-On Tutorial

Adamo Tonete, Vinicius Grippa, Percona

MongoDB® is the most popular NoSQL database in the world and the 5th in general. Thanks to its flexibility storing data, horizontal scaling and high availability features MongoDB has its presence in the hall of fame.

We will discuss how data is stored on MongoDB and the internal aspects of the high availability and how to scale it. Finally, one look closer to the future and how to deploy MongoDB on Kubernetes.

It is expected that all participants (developer or operations personnel) will leave more knowledgeable on how MongoDB works, main use cases and how to setup a standalone, replicaset and how to use Percona MongoDB Operator for Kubernetes.

Attendees will undertake practice sessions on:

- Basic concepts of MongoDB
 - » Installation, best practices configuration, user and roles
- How the replicaset works
 - » What is a heartbeat, how does an election works, hidden instances, read preference
- MongoDB on Kubernetes (how to deploy, self-healing, auto-escalating)
 - » Create a replicaset with Percona Kubernetes Operator for Percona Server for MongoDB in 3 minutes with no manual intervention to setup the replication

Accelerating Application Development with Amazon Aurora

Kathy Gibbs, John Nguyen, Amazon Web Services

In this a hands-on tutorial, you'll learn how to leverage the unique features of Amazon Aurora to build faster, more scalable database applications optimized for the cloud. We discuss architectural best practices and features designed to help you develop applications faster and reach the widest possible audience, including Aurora Serverless, read replica auto scaling, cross-region replicas, backtracking, fast database cloning, and Performance Insights. You'll understand how to best take advantage of the Aurora platform's capabilities to effectively accelerate application development.

Please note: this is a half day tutorial that will be repeated in morning and afternoon sessions.

**CONFERENCE DAY ONE
TUESDAY, OCTOBER 1**

Our Journey to Better MySQL Availability Using Global Transaction IDs, ProxySQL and Consul

Stephane Combaudon, Yuriy Olshanetsky, Rakuten Intelligence

Data is core to our business at Rakuten Intelligence. And while databases not being available can hurt us badly, our system was not initially designed with MySQL high availability in mind. Each MySQL master switch involved stopping and restarting applications, which at some point was no longer a valid option.

No existing HA solution worked well for us, therefore we focused on implementing the glue between existing bricks, the main ones being: ProxySQL (connection management between the applications and MySQL), Consul (service discovery) and MySQL GTIDs (easy tracking of transactions across servers).

We will share our failures and how we incrementally improved our system to reach our availability goals.

More Than a Query Language: SQL in the 21st Century

Markus Winand, winand.at

Did you know the purely relational dogma of SQL was already abandoned in 1999? The last SQL standard that was limited to the relational idea was SQL-92. From SQL:1999 onwards, the SQL language was extended with non-relational operations and non-relational data structures. As much as this move was discussed at that time, it took decades until database vendors caught up with this idiomatic change. Many SQL users haven't heard of it until today.

This talk provides the big picture on the evolution of the SQL standard and introduces some selected modern SQL features by example. You will see that SQL has changed as much as our requirements have changed over the past decades.

MySQL 8.0 Performance: Scalability & Benchmarks

Dimitri Kravtchuk, Oracle

This session will be interesting to everyone looking for the latest news about MySQL 8.0 Performance:

Since MySQL 8.0 we moved to “continuous release” model so with every update many new improvements are delivered but how does it also improve MySQL 8.0 Performance ? ;-)

The latest benchmark results obtained with MySQL 8.0 will be in center of the talk because every benchmark workload for MySQL is a headache and “problem to resolve” and each resolved problem is a potential gain in your production !

Many important internal design changes are coming with MySQL 8.0

- how to bring them all in action most efficiently ?
- what kind of trade-offs to expect, what is already good, and what is not yet ?
- how well MySQL 8.0 is able to use the latest new HW and how can it run efficiently on small boxes ?
- could you really speed-up your IO by deploying your data on the latest flash storage ?

These and many other questions are answered during this talk + proven by benchmark results... and as usual, some surprises to expect ;-))

Vitess: Running Sharded MySQL on Kubernetes

Sugu Sougoumarane, PlanetScale

Vitess has continued to evolve into a massively scalable sharded solution for the cloud. It's is now used for storing core business data for companies like Slack, Square, JD.com, and many others.

This session will cover the high-level features of Vitess with a focus on what makes it cloud-native.

We'll conclude with a demo of the powerful materialized views feature that most sharded systems have yet to solve.

Amazing Sandboxes with dbdeployer

Giuseppe Maxia, dbdeployer

You probably know that you can deploy a MySQL database or a group of databases in replication in your laptop using dbdeployer.

But unless you have been paying close attention, there are a few brilliancies that may have escaped you. This session will show the basic usage of dbdeployer, before delving into the advanced operations, such as: importing running databases into a sandbox, replicating on-the-fly between single sandboxes, or between groups, or NDB clusters. You can also clone sandboxes with a single command (sometimes with just half a command) and run commands in several servers at once. Want to run a query on all slaves? You will see that! How about on all masters? You've got it!

Want to change a sandbox server configuration temporarily or permanently? You will see that also!

Do you want to make your own commands with fine tuning options depending on database flavor and version? dbdeployer sandboxes are built for this purpose: you just need to see how to use the force that has been waiting for you.

Multi-Document Transactions Through MongoDB Shell

Aayushi Mangal, Percona

Transactions are a well known and widely used feature of RDBMS systems, but MongoDB is a document-oriented non-RDBMS database. Every single document is atomic in MongoDB,. It is widely known for its simplicity of importing data, and the sharding of data across many servers. This talk includes transactions in detail and internals of transactions oplogs.

Key points covered:

- Transactions
- Prerequisites
- writeConflicts
- Considerations and best practices
- Hands-on through MongoDB shell

Tarantool: A NoSQL Database with SQL **Pavel Lapaev, Tarantool**

We will present a short history and overview of Tarantool, an open source in-memory DBMS. We will explain why it is cool to have a DBMS in the same address space as your application server, why Tarantool is in fact single threaded and other main features. Then, we'll explain, why a no-SQL DBMS supports SQL.

Percona XtraDB Cluster 8.0 (PXC 8.0) **Krunal Bauskar, Percona**

Percona XtraDB Cluster 8.0 (PXC-8.0) is the latest addition to PXC family. Starting with MySQL/PS-8.0, upstream have made a lot of significant changes including atomic DDL, replication channel, locking algorithm changes, etc....

During this session, we will explore:

- how these changes affect PXC-8.0
- what has changed in PXC-8.0
- new/deprecated features
- important bugs and more.

If you are already a PXC user or planning to consider it, attend this session to find out more about PXC-8.0.

What's New in PMM 2.0 **Michael Coburn, Rachel Pescador, Percona**

PMM 2 represents a significant advance in terms of monitoring for Open Source Databases. Come to this session in order to learn about the following:

- Query Analytics improvements
- Architectural changes
- API and GUI configuration

MySQL on Google Cloud: The Good, The Bad, and The Ugly

Jeremy Cole, Shopify, Inc.

Running MySQL in the cloud isn't magic - it can be brilliant, but it can also be a real challenge. We'll learn about some of the big wins that can be had in the cloud (such as elasticity and easy provisioning). But, we'll talk about the dark underbelly as well - and face some of the challenges to run a realistic MySQL installation in the cloud (errrr... performance?).

MongoDB Data Security: Custom Rules and Views

Adamo Tonete, Percona

In this presentation, we will discuss how to create custom rules when the default rules are not enough for the application. Have you needed to give a more permissive rule to a user just because of this user wanted to run a specific command? Also, we will discuss how to use view for hiding fields from users when we don't want them to read all the collection. If you have concerns about security come to this talk.

Why PostgreSQL is Becoming a Migration Target in Large Enterprises

Jobin Augustine, Percona

1. Integration with Other RDBMS systems, Oracle, MSSQL, and MySQL
2. Integration with Hadoop and NoSQL databases
3. Compatibility layer: Example of OraFCE
4. Migration tools: Example of Ora2Pg
5. How PostgreSQL is suitable for Agile development and Knowledge Portability.
6. PostgreSQL derived projects and freedom of choice.

How to Protect PII Data in MySQL While Allowing Customers to Access the Database

Alexander Rubin, VirtualHealth

Protecting PII data is one of the biggest challenges of security minded DBA. But when you add a requirement to allow customers to have a direct access to a database replica it will make a life of a DBA nightmare. At VirtualHealth we deals with medical data. In this talk I will share an experience of how to implement customer authentication, authorization and data isolation in MySQL. I will cover such topics as:

- Authentication of external users: how to create a secure connection to MySQL and manage users (LDAP, external auth, etc)
- Authorization and data isolation
- Auditing the logins and login attempts
- Encryption at rest, field level encryption and key rotation
- Data masking and data de-identification (for dev systems)

MyRocks and RocksDB Advanced Features and Performance

Yoshinori Matsunobu, Facebook

We have been running MyRocks in Facebook's largest database (UDB) and Facebook Messenger backends for a while. We have made significant improvements in both MyRocks and RocksDB, especially from performance, efficiency, reliability and operations point of view.

In this session, the speaker will go through several important features and performance improvements.

New Indexing and Aggregation Pipeline Capabilities in MongoDB 4.2

Antonios Giannopoulos, ObjectRocket by Rackspace

MongoDB 4.2 comes GA soon delivering some amazing new features on multiple areas. In this talk, we will focus on the new capabilities of the aggregation framework. We are going to cover the new operators and expressions. At the same time, we will explore how updates commands can now use the aggregation framework operators. We are also going to present aggregation framework improvements focusing on the on-demand materialized views. Finally, we are going to explore the wildcard indexes introduced in MongoDB 4.2 and how they change the way we design documents and build queries/aggregations. We will also make a reference to the new index build system.

Comparison of Kubernetes Operators for MySQL

Mykola Marzhan, Percona

Running databases in Kubernetes attracts a lot of attention today. Orchestration of MySQL on Kubernetes is no way a straightforward process. There are several good MySQL based solutions in the open-source world, made by Oracle, Presslabs, and Percona. Having a common base, they differ in self-healing capabilities, multi-master and backup/restore support, etc. So let's make a fair comparison to figure out the pros and cons of their current state.

Pg_catalog Unrevealed! That Part of PostgreSQL You Are Probably Underusing

Boriss Mejias, 2ndQuadrant

PostgreSQL users want to know how the system is performing. This is valid for DBAs, system administrators, and developers. They want to know if the system needs some improvements, but also to generate the occasional report requested by managers. Standard monitoring tools will monitor the CPU, RAM and I/O consumption, but they won't be able to tell you if indexes are being used, if tables are bloated, what is the lag of replication, what are the status of sessions connected, and many other interesting and useful things from PostgreSQL.

This talk is intended to review the possibilities offered by the PostgreSQL catalog. Plenty of information is available there but it is unknown to many users, and therefore unused. We will see how to exploit the catalog, how to send the information to other monitoring tools, and how the tables in pg_catalog are fundamentally interconnected to other topics such as performance, replication, MVCC, security, the universe and everything.

Enhancing MySQL Security

Vinicius Grippa, Percona

Security is always a challenge when it comes to data protection but regulations like GDPR and HIPAA brings a new layer on top of it with severe fines for breaches and rules more restrictive to access and manipulate personal data. Join us in this presentation to check security best practices, traditional and new features available for MySQL 8.

• On this talk, DBAs and sysadmins will walkthrough since the beginning in the application where data is born until the end it is archived on the backup. The topics are presented below:

- SO/Cloud security
- SSL
- Password Management
- TDE (Transparent Data Encryption)
- Audit Plugin
- MySQL 8 features
- New caching_sha2_password
- Roles
- Password Management
- FIPS mode
- Using Xtrabackup to encrypt the backups

We will share our experience of working with 2000 support customers and help the audience to become familiar with all the security concepts and methods, and give you the necessary knowledge to apply to your environment.

PCI/DSS Compliance with MySQL: 2019 Edition

Carlos Tutte, Percona

In the past decade there was a boom in online transactions, and because of this criminals moved online to exploit billions of insecure consumer information records. Payment Card Industry Data Security Standard (PCI/DSS) dictates norms about storing, processing and transmitting cardholder data and originated in an effort to reduce and prevent credit card fraud. Discover which tools/plugins/settings you need to use to comply with PCI/DSS when using MySQL, including:

- Protect cardholder data from unauthorised use
- Enforce strong controls around privileged users and data access
- Implement centralized, automated role-based access control, authorization, and authentication
- Provide system and database auditing, and database activity monitoring
- Encryption in MySQL: TDR, LUKS, AES, DES, asymmetric key

Attendees will gain a solid grasp of the possibilities MySQL offers to someone looking to become PCI/DSS certified, as well as the limitations it comes with.

What's New in Percona Server for MongoDB

Akira Kugonane, Adamo Tonete, Percona

News on the free-and-open enterprise features in PSMDB (v4.0+) and the recent v4.2 release.

Data-at-rest Encryption requires key management (no point storing the keys with the data!). We'll demonstrate how v4.0.10-5 of Percona's fork uses integration with Hashicorp Vault to provide this.

Percona Server for MongoDB 4.2.0 has been released merging all of MongoDB's upstream community edition code. The headline features of distributed transactions and update-capable aggregation pipelines are in (see other talks!). In this presentation we'll share some other important changes that all MongoDB DBAs should be aware of.

- A new throttle mechanism that works on primary nodes;
- A merging of foreground and background index build methods
- The eight-year wait to be able to list all client ops (active or idle) is finally over.
- Wildcard Indexes

Don't miss this talk!

Lessons From Building Automation for a Large Distributed Database

Ameet Kotian, Leigh Johnson, Slack Inc

Automation is a critical investment when scaling infrastructure operations. If you delay automating, your company's growth will end up supported by a mountain of human toil.

Everyday operations like provisioning, host replacements start of as steps in a runbook. Later, you might code-ify these processes with scripts. At some point, there is enough complexity and requirements for your automation that it evolves into a piece of full-fledged production software.

The automation of stateful services, like databases, in particular comes with a special set of problems, including:

- Risk of data loss
- Strict SLOs for availability
- Longer provisioning times due to data restore process and migrations

The automation needs additional level of safety built-in to minimize data loss risk. This is a common problem without a common, universal solution. We needed to figure out an answer to this problem while scaling out a stateful MySQL database fleet to thousands of hosts.

In this talk, we will outline the natural evolution of automation software. We believe, at scale automation software, has the same sets of rigorous requirements as any other production software but there are easily identifiable patterns common to automation software - things like - a scheduler/state-machine to execute a series of steps, resiliency features such as timeout and retries, method of failure handling and escalation, audit features such as logging and actions history. Finally some sort of interface clients like web-ui or command line.

Once such an automation framework is available, we can build self-healing infrastructure on top it by feeding it reliable and robust failure signals.

ClickHouse for Time-Series Real-time Analytics

Alexander Zaitsev, Altinity

ClickHouse is a very fast general purpose open source DBMS. Starting from web analytics use case several years ago, it quickly develops to other areas, including time series applications where ClickHouse manages to outperform dedicated time series solutions. In this talk I will explain different approaches of using ClickHouse for time series, highlight ClickHouse features that help to design an efficient schema, present some benchmarks, and walk through real life optimisation techniques.

Improving Enterprise HA and Disaster Recovery Solutions

Marco Tusa, Percona

Who nowadays does not know of HA solutions based on MySQL/Maria/Percona cluster (Galera)?

Even so, the rollout of this very powerful solution is often not correctly designed, especially when trying to solve not only local HA, but also to meet Disaster Recovery scenarios.

During the many experiences since my first MySQL+Galera installation, I have faced many misconceptions when customers have needed to design solid/trustable/efficient DR solutions.

Foremost is the lack of consideration of what can be really achieved without breaking the laws of physics.

This presentation wants to clarify once for all WHY some of the current misconception are ...just that, "misconceptions".

At the same time, I will illustrate this talk using an architecture solution, that we have identified as the real answer for DR scenarios and that we have successfully implemented many times.

I will start from customer requirements, and show how we have guided our customers through the available options to finally identify the correct solution.

We will discuss the right way to analyze the platform, and how to finally setup a MySQL/Maria/Percona cluster.

The presentation will also answer the question on how setup multiple DR sites.

Join Heterogeneous Databases Using PostgreSQL Foreign Data Wrappers

Ibrar Ahmed, Percona

PostgreSQL provides a way to communicate with external data sources. This could be another PostgreSQL instance or any other database. The other database might be a relational database such as Clickhouse, MySQL or Oracle; or any NoSQL database such as MongoDB or Hadoop. To achieve this, PostgreSQL implements ISO Standard call SQL-MED in the form of Foreign Data Wrappers (FDW). This presentation will explain in detail how PostgreSQL FDWs work. It will include a detailed explanation of simple features and will introduce more advanced features that were added in recent versions of PostgreSQL. Examples of these would be to show how aggregate pushdown and join pushdown work in PostgreSQL.

The talk will include working examples of these advanced features and demonstrating their use with different databases. These examples show how data from different database flavors can be used by PostgreSQL, including those from heterogeneous relational databases, and showing NoSQL joins.

Backing up Wikipedia Databases

Jaime Crespo, Manuel Arostegui, Wikipedia

Imagine a world where all Wikipedia articles disappear due to a human error or software bug. Sounds unreal? According to some estimations, it would take an excess of hundreds of million person-hours to be written again. To prevent that scenario from ever happening, our SRE team at Wikimedia recently refactored the relational database recovery system.

In this session, we will discuss how we backup 550TB of MariaDB data without impacting the 15 billion page views per month we get. We will cover what were our initial plans to replace the old infrastructure, how we achieved recovering 2TB databases in less than 30 minutes while maintaining per-table granularity, as well as the different types of backups we implemented. Lastly, we will talk about lessons learned, what went well, how our original plans changed and future work.

MySQL InnoDB Cluster: Advanced Configuration & Operation

Frédéric Descamps, Pedro Gomes, Oracle

MySQL InnoDB Cluster is a very easy HA solution to deploy.

However it's also a very customizable solution able to respond to most needs. During this session I will give an overview of settings that you may tune like those related to quorum lost, level of consistency, but also some you may not know like how to change recovery system, effect of increasing the event horizon. We will also discuss about maintenance operations like how to stream large transactions, how to deal with DDL in multi-primary environments...

InnoDB Management and Scalability Improvements in MySQL 8.00

Sunny Bains, Oracle

At the beginning there was one general tablespace... Then we added a tablespace per table... What have we designed and delivered in MySQL 8.0 to manage even better your tablespaces?

Come to this session to learn everything you need about InnoDB tablespaces.

Deep Dive on Amazon Aurora **Yoav Eilat, Amazon Web Services**

Amazon Aurora is a MySQL- and PostgreSQL-compatible relational database designed from the ground up to benefit from the cloud. Aurora offers a novel architecture that moves the monolithic database stack to a service-oriented architecture, pushing the lowest layers of the database into a distributed, multi-tenant, log-structured storage service. We'll cover key Aurora capabilities such as low latency replicas, instant crash recovery, in-place rewind, copy-on-write cloning, zero-downtime patching, multi-master, and serverless, and look under the hood to explain how they work.

High Availability and Automatic Failover in PostgreSQL

Avinash Vallarapu, Percona

PostgreSQL is an advanced open source database that is completely community driven. Continuous development and performance improvements while making it feature rich gave PostgreSQL a great market penetration. When a database server is deployed in production, it is important to achieve several 9's of availability. It starts with a replication setup that is stable and some tool that could manage a failover transparent to application. This talk starts with a list of replication types in PostgreSQL followed by a detailed understanding of some of the widely used open source tools for high availability and automatic failover in PostgreSQL.

The following topics will be discussed during this talk:

1. How the implementation of HA differs for each type of replication available in PostgreSQL
2. How to combine HAProxy with etcd plus a detailed explanation of the RAFT Algorithm for HA using Patroni.
3. What is stolon, pgpool-II and repmgr and how can they be helpful.
4. Some other tools admins may wish to explore.
5. What are the tools and solutions that help you achieve automatic failover in AWS and other cloud-based environments.
6. How to avoid huge data loss during failovers.

VictoriaMetrics: Why and How We Built Fast and Scalable Open Source Time Series Database

Roman Khavronenko, Aliaksandr Valialkin, VictoriaMetrics

How did we get the idea of creating new time series database?

The talk describes typical evolution of application monitoring during the last few years: it starts with zero monitoring, then logging is added, then the observability is improved with app-specific metrics, altering and graphing. The evolution meets various stumbling blocks - scalability, performance, storage space usage. They may result in creating new TSDB such as VictoriaMetrics.

From Scheduled Downtime to Self-Healing in Less Than a Year

Károly Nagy, Salesforce

Infrastructure automation is not easy, especially for stateful services like MySQL (or any other database for that matter). It goes way beyond the capabilities of Ansible, Chef, SaltStack or other similar tools. In this session I'm going to show you how we went from fully manual operations to a self-healing system in less than a year at Salesforce. Having done this at several companies already I've seen the common mistakes that can break your system and make your well intended scheduler/scripts/orchestrator a ticking bomb. I will share how to avoid these problems and build a robust and scalable automation framework that's been battle tested at companies such as Booking.com and Dropbox.

We will cover:

- Tool comparison
- Centralised vs decentralised system
- Concurrency handling
- Best practices and anti-patterns

MySQL 8.0.18: Latest Updates

Norvald Ryeng, Oracle

MySQL 8.0.18 is almost here! In this talk we'll focus on two of the new features: hash joins and EXPLAIN ANALYZE. This talk will cover how these features work, both on the surface and on the inside, and how you can use them to improve your queries and make them go faster.

Both features are the result of major refactoring of how the MySQL executor works. In addition to explaining and demonstrating the features themselves, we'll look at how the investment in a new iterator based executor prepares MySQL for a future with faster queries, greater plan flexibility and even more SQL features.

Tracing and Profiling MySQL

Valerii Kravchuk, MariaDB Corporation

While troubleshooting MySQL performance problems it is important to find out where CPU time is spent inside mysqld process. The process of investigation should have as small influence as possible on the server we try to troubleshoot.

Performance_schema introduced in MySQL 5.5 is supposed to provide detailed enough instrumentation for most cases. But it comes with a cost, require careful sizing of performance counters, and the process of instrumenting the code is not yet complete even for MySQL itself, to say nothing about 3rd party storage engines, plugins and libraries like Galera.

This is when profilers come handy. Poor Man's Profiler (PMP, aka pt-pmp) and perf can be easily used while studying MySQL performance problems. Basic usage steps are presented and several typical use cases are discussed.

Tracing system calls (that is needed for MySQL troubleshooting sometimes) is usually performed with strace. But recent Linux kernels with eBPF support allows to trace these and more with less performance impact and collect more details. Brief introduction to bpftrace is presented.

ClickHouse Features to Blow Your Mind

Aleksei Milovidov, Yandex

I will present a subset of the most notable ClickHouse features over the last half of year:

- data skipping indices, including full text indices (with performance evaluation and insights on implementation)
- custom compression codecs for time series data;
- HDFS and Parquet integration;
- fuzzy string search (it is really fast fuzzy string search); multiple substring matching;
- sampling profiler on the query level;
- z-curve indexing;
- table and columns TTL;

OpenCorporates: Providing Transparency for the Public Benefit

Ben Tullis, OpenCorporates

OpenCorporates is one of the largest open databases made available for the public benefit in the world. It hosts structured data on over 100 million companies and is used regularly by journalists, investigators, governments, even banks.

MySQL continues to be a crucial part of our technology stack and the size of the primary database has increased by around 800% in the past four years.

This talk describes our journey and the challenges, pitfalls and crucial business decisions along the way which led us to our current implementation using Percona XtraDB Cluster and ProxySQL on top of ZFS.

TUESDAY

[RETURN TO TUESDAY'S TIMETABLE](#)

CONFERENCE DAY TWO
WEDNESDAY, OCTOBER 2

Parted Ways with Partitioning? It's Time to Reconsider...

Feike Steenbergen, Timescale

How PostgreSQL Revived Partitioning to Improve Database Performance

Partitioning has improved drastically in PostgreSQL in the last 3 releases.

Whereas a few years ago partitioning would only be considered to solve specific pain points, it is now in such a state that you can use the power of partitioning in almost any database design to improve performance, to aid in data retention policies, and to have fewer operational headaches.

In this talk, we'll briefly go through the history of partitioning, and then we will take a deep dive into what this means for your next schema (re)design. More specifically, we will discuss:

- Exploring design patterns
- Validating your database and schema designs
- What to look for in your explain plans
- How to ensure that partitioning pruning is done at the correct time

The way we will answer these questions is by using some of the most useful performance measurement tools available in PostgreSQL: EXPLAIN and pg_stat_statements.

We'll use a publicly available data set and see how both the performance and the query plans have changed in the past releases of PostgreSQL.

Percona Server for MySQL: What It Is and How It Is Done

**Satya Bodapati, Yura Sorokin,
Percona**

This talk will explain what exactly Percona Server is and how it is similar to / different from Oracle MySQL Server. It will also give you a brief description of some unique features it already has and may have in the near future.

In addition, it will reveal a few secrets about what Percona Dev Team does on regular basis to keep up-to-date with Oracle MySQL, to constantly improve code quality and to fight with bugs.

Billion Goods in Few Categories: How Histograms Save a Life?

Sveta Smirnova, Percona

We store data with an intention to use it: search, retrieve, group, sort... To do it effectively the MySQL Optimizer uses index statistics when compiles the query execution plan. This approach works excellently unless your data distribution is not even.

Last year I worked on several tickets where data follow the same pattern: millions of popular products fit into a couple of categories and rest used the rest. We had a hard time to find a solution for retrieving goods fast. We offered workarounds for version 5.7. However new MariaDB and MySQL 8.0 feature: histograms, - would work better, cleaner and faster. The idea of the talk was born. Of course, histograms are not a panacea and do not help in all situations.

I will discuss:

- how index statistics physically stored by the storage engine
- which data exchanged with the Optimizer
- why it is not enough to make correct index choice
- when histograms can help and when they cannot
- differences between MySQL and MariaDB histograms

MySQL Has Gone Away: An In-Depth Look at the MySQL Networking Implementation

Iwo Panowicz, Percona

We'll take a look at the implementation of the network layer in MySQL 8.

- What is the purpose of Virtual I/O Library?
- How do packets move from a socket to Storage Engines?
- How network related algorithms work in MySQL?
- poll(2), select(2), ppoll(2), epoll(2) - what you can find in MySQL source code?
- What's the real role of MySQL network variables (like net_read_timeout, net_write_timeout, back_log etc.)? Why usually are defaults the best choice?
- How not closing connections by applications affects MySQL?
- When to tune MySQL network layer and how it affects MySQL?
- MySQL has gone away from a low-level perspective.
- SSL/TLS, is it slower? And why?
- When to tune Operating System network stack?

How a Modern Database Gets Your Data Fast: MariaDB Query Optimizer

Vicențiu Ciorbaru, MariaDB Foundation

A modern database query optimizer needs to be able to handle a lot of different workloads, especially when it comes to general-purpose databases. Most of the time, it does a good job; however, certain patterns are hard to predict. With so many moving parts, it's hard for the average DBA or database developer to come up with a good explanation for why the optimizer chooses certain query plans.

Vicențiu Ciorbaru dives deep into the optimizer, covering strategies for dynamic query rewriting, including derived-table merging and condition push down and moving HAVING clauses into WHERE; temporary table reuse; join optimizations via IN/EXISTS subqueries; condition selectivity estimations via pre-collected histograms; and improved analytics via Window functions and CTEs. Armed with this extra information, you'll be able to better tune problematic queries and better model a database schema, especially for large-scale datasets.

The talk will focus on MariaDB's optimizer point of view; however, he examines other database systems where it's relevant.

Maintenance for MongoDB Replica Sets

Igor Donchovski, Pythian

Redundancy and high availability are the basis for all production deployments.

With MongoDB high availability is achieved with replica sets which provides automatic failover in case the Primary goes down. In this session we will review multiple maintenance scenarios that will include the proper steps for keeping the high availability while we perform the maintenance steps without causing downtime.

This session will cover Database upgrades, OS server patching, Hardware upgrades, Network maintenance and more.

- How MongoDB HA works
- Replica sets components/deployment typologies
- Database upgrades
- System patching/upgrade
- Network maintenance
- Add/Remove members to the replica set
- Reconfiguring replica set members
- Building indexes
- Backups and restores

Percona Backup for MongoDB **Akira Kurogane, Percona**

Percona's open-source solution saves consistent sharded cluster backups to S3 (or S3-compatible) object stores or locally-mounted filesystem. It also provides a single-command action to restores those in-place. It is compatible with MongoDB® Community edition as well as Percona Server for MongoDB,

Please join to hear about use cases and the simple topology used for PBM agents.

Databases, Crypto & Decentralization

Caleb James DeLisle, CJDNS

The first major revolution in computing was lead by the database, and from the mainframe era to the rise of web2.0, the venerable database has been at the heart of every information system. However, new technologies such as federated social networking and blockchain threaten to sideline the database in favor of the protocol, even calling into question it's future relevance.

I will argue that in the era of decentralization, the database is more relevant than ever, and that the next revolution in computing will be the general purpose decentralized database, where all manor of decentralized apps from federated social networking to cryptocurrency can coexist as schemas in a DBMS.

ClickHouse Materialized Views: A Secret Weapon for High Performance Analytics

Robert Hodges, Altinity

Materialized views are a powerful addition to data warehouses that automatically transform data into new forms and improve query performance by orders of magnitude. ClickHouse has a particularly good implementation. This talk introduces ClickHouse materialized views using standard examples that should be in the toolkit of any data warehouse designer. We start with simple aggregation and extend to time series use cases like the 'last point problem,' reclustering data for better performance, and even pulling data from Kafka queues. Through these concrete examples you'll understand the benefits of materialized views as well as the plumbing that makes them work in large, distributed data warehouses.

Fortify Your MySQL Data Security in AWS Using ProxySQL Firewall

Marco Tusa, Percona

In Percona we care about you, and we are constantly looking for the right solution who can answer your needs.

Securing our data is a complex topic. We can build a very strong protection around our data, but nothing will prevent the one WHO could potentially access it to compromise the data integrity or to expose it.

This because we either under estimate the control we can or should impose, or because we think to do not have the tools to perform such control.

Nowadays to be able to control and manage what can access our data is a must, while how to do with standard tools it is a nightmare.

The presentation will guide you in a journey, there you will discover how implementing a quite robust protection, more than what you thought was possible.

Even more, it is possible and your performances will even improve. Cool right?

We will discuss:

- Access using not standard port
- Implement selective query access
- Define accessibility by location/ip/id
- Reduce to minimum cost of filtering
- Automate the query discovery

Percona Distribution for PostgreSQL

Ibrar Ahmed, Percona

Percona has announced the release of Percona Distribution for PostgreSQL. Alongside the PostgreSQL distribution, our package incorporates other tools that we feel native PostgreSQL lacks: pgBackrest, a better alternative to PostgreSQL pg_basebackup, and Patroni to support an HA solution.

Percona provides the rpm and Debian packages for PostgreSQL and these tools. I will discuss how to enable a repository to use the Percona distribution, and how some of the extensions work.

Currently, we are at the beta stage, stay tuned for the GA!

MySQL Shell: The Best DBA Tool? Extend the Shell With the New Extension Infrastructure

Frédéric Descamps, Oracle

MySQL Shell 8.0.17 has been released with a new extension framework. Join this session to discover how to write your extension. I will also give you an overview of several useful extensions, how to find them and how to install them.

Finally I will show you how to create your own.

Data Protection and OSS in the Age of GDPR

Cristina DeLisle, DPO XWiki SAS

Privacy is becoming central in shaping the future of tech and the data protection legislation has contributed significantly to making this happen. Privacy by default and design are core principles of how software should be envisioned. The GDPR that came into the spotlight last year has a strong case to become a standard even outside European borders, influencing the way we protect personal data. However its impact might be, its implementation is still in its infancy. OSS has found itself facing the situation and one aspect which is particularly interesting on the tech side is how to incorporate the principles of privacy by default and design into the software that we build.

This talk is an overview of how the GDPR has impacted OSS, how to apply the model of data controller and data processor on the contexts of various communities and infrastructure providers. Also, it will tackle aspects regarding what do we mean by privacy by default and by design, how could we envision them applied in our OSS. It will bring examples from which we might find something interesting to learn from, regardless if we are looking at them as mistakes, best practices or just ways of doing things.

What's New On Sharding in MongoDB 4.2

Antonios Gianopoulos,
ObjectRocket by Rackspace

MongoDB 4.2 comes GA soon delivering some amazing new features on multiple areas. In this talk, we will focus on changes related to sharded clusters. We are going to cover distributed transactions & mutable shard keys providing examples that will reveal the internals of those new features. We will provide best practices around the new sharding features and we will cover other minor changes related to it.

My First 90 Days with Vitess **Morgan Tocker, PlanetScale**

I have worked with MySQL for several years, but my adventures with Vitess have only just started. With this experience still fresh in my mind, I want to share with other MySQL users looking to walk down the same path what I found most interesting to discover :-)

This talk will start with the basic architecture of Vitess. From there, I will go over some of the basic concepts and terminology that I found useful to learn such as Topos, Keyspaces, VSchema, and Tablets. Finally, we will talk about MySQL compatibility. This includes both what schemas and SQL you can expect to be supported.

TO BE CONFIRMED

Running PMM in Production at Tessi

Valentin Traën, Jonathan Gourdan, Tessi & Michael Coburn, Percona

Have you struggled with problem identification using MySQL?

The sooner problems are identified means the faster the environment is stabilized.

Come listen to how Tessi Document Services uses PMM + MySQL InnoDB Cluster to visualize cluster performance and to identify & analyze slow queries.

We'll also show some of the dashboard and custom scripting improvements made to extend PMM.

Handling Transaction ID Wraparound in PostgreSQL

Avinash Vallarapu, Percona

A transaction ID in PostgreSQL is a 32 bit unsigned integer. 2.1 billion transactions in the past are visible and 2.1 billion in the future are not visible. What it means is, a tuple inserted by a transaction ID : n is not visible to a transaction id : $(2.1 \text{ billion} + n + 1)$. This is called a transaction ID wraparound. However, there is a mechanism in PostgreSQL that avoids this from happening. But, there can be a situation when it could go so near to that age where downtime is needed to perform a maintenance task. In this talk, the audience would get introduced to the internals of transaction ID's with a detailed understanding of transaction ID wraparound. And then understand how this can be avoided through parameter tuning and routine monitoring and maintenance. And finally move to the procedure involved in repairing the damage when a database is no more available for writes.

This talk covers the following topics:

- Transaction IDs in PostgreSQL
- How does a wraparound happen?
- How is it avoided?
- What happens before the age of 10 million and 1 million of 2.1 billion?
- Should template databases be vacuumed?
- How to monitor and how to tune?
- How to repair damage?

Securing MySQL with Vault and Table Level Encryption

Kevin Markwardt, Pythian

With all of the leaks of personal information it is very important to make sure that your data is secure and protected. In this talk I will go into and demonstrate how to secure your MySQL user accounts with Hashicorp Vault, to do MySQL user account rotation and expiration. You can then integrate it with LDAP. There are a lot of users that do not use command line interface that may want to access these environment. This can all be done with Vaults built in GUI interface. Finally I will demonstrate how to encrypt your tables and store the key in Vault by base encoding your key and placing it into a KV store, or by just using the Percona keyring_vault plugin.

Benchmarking Should Never Be Optional

Art van Scheppingen, Messagebird

Benchmarking is evaluating your database system or environment by comparing it with a set baseline. In this session you will learn why benchmarking any (database) system is not optional and should never be skipped. Benchmarking takes time and often, as projects are often rushed into production, skipped due to time constraints and this means you will be flying blind as you don't know what your limits are. Also a common misconception is that benchmarking is only to be used for finding the optimal performance of your database, while it can give you a whole additional array of insights on your entire environment: resource planning, predict behaviour and even give insight in the Mean Time Till Repair (MTTR).

Leveraging benchmarks for resource planning will give you insight into how much capacity you need or when you are expected to scale or upgrade. Benchmarks can also be utilized to put an application or architecture under a lot of stress. This can predict how the application or entire architecture will behave when it's resources near exhaustion and this can help you to identify issues early on when this would otherwise happen when it hits production.

Finally you can also see the process of optimizing your Mean Time Till Repair as a benchmark exercise, which can be part of your disaster recovery plan.

10 Common Mistakes (Java) Developers Make when Writing SQL

Charly Batista, Percona

It's easy for Java developers (and users of other OO languages) to mix object-oriented thinking and imperative thinking. But when it comes to writing SQL the nightmare begins! Firstly because SQL is a declarative language and it has nothing to do with either OO or imperative thinking and as for one point it makes it relatively easy to express a condition in SQL it is not so easy to express it optimally and even worse to translate it to the OO paradigm. For another point, they need to think in terms of set and relational algebra, even if unconsciously!

In this talk, we'll see the most common mistakes that OO, and in particular Java, developers make when writing SQL code and how we can avoid them.

Top 10 Mistakes When Migrating From Oracle to PostgreSQL

Jim Mlodgenski, Amazon Web Services

As more and more people are moving to PostgreSQL from Oracle, a pattern of mistakes is emerging. They can be caused by the tools being used or just by not understanding how PostgreSQL is different than Oracle. In this talk, we will discuss the top mistakes people generally make when moving to PostgreSQL from Oracle and what the correct course of action.

Large Scale Deployment of SSL/TLS For MySQL

Daniël van Eeden

Deploying SSL/TLS with MySQL at Booking.com on thousands of servers is not without issues.

In this session I'll tell you what steps we took, what problems we hit, and how we improved various parts of the MySQL ecosystem while doing so.

To start we go over the basics: Which TLS settings are there in MySQL and MariaDB and how does this differ from HTTPS as used in browsers. And why do we want TLS in the first place? Is TLS and SSL the same thing?

The first set of problems is inside MySQL: YaSSL vs. OpenSSL, verification issues and reloading of certificates.

The second set of problems is inside Connectors: I'll touch on DBD::mysql (Perl), Go-MySQL-Driver, libmysqlclient (C)

Not all connectors have the same options and defaults. I'll go into TLSv1.2 support.

The third set of problems is tools: Using the `require_secure_transport` option caused issues with Percona Toolkit and Orchestrator.

I'll also cover: RSA v.s EC, security issues I found and how I wrote a Proxy for MySQL

Protecting Your Secrets

John Boero, Hashicorp

When we talk about secret management what we're really talking about is managing a set of different credentials. What we mean when we talk about these credentials is anything that might grant you authentication to a system. Some examples of this might be usernames and passwords, it might be things like database credentials. This falls into the realm of secrets, and these are things we want to manage carefully. We want to understand who has access to them, we want to understand who's been using these things, and in the case of most of these, we want some story around how we can periodically rotate these. In this session, John will present his recommendations on how to protect your secrets and avoid secret sprawl.

Strength in Numbers:

Slack's Database Architecture

Josh Varner, Guido Iaquinti, Slack

Traditionally, database reliability has focused heavily on the stability of a single server or a small number of servers. At Slack, however, we've built a database architecture that instead focuses on an approach of strength in numbers. We've designed this architecture to allow failures to happen and to limit the effects of those failures instead of attempting to eliminate them entirely. We've also built this architecture from the ground up to be cloud-first, and to use the nuances of running databases in the cloud to our advantage.

In this talk, we will review the architecture we've built, specifically focusing on our Vitess infrastructure. Vitess currently handles nearly half of Slack's database traffic. We'll talk about how we've built this system to withstand common failures and continue to perform at our scale.

MySQL 8.0: The New Replication Features

Luis Soares, Oracle

MySQL 8.0 replication reduces operation costs and minimizes administration overhead. MySQL DBAs and DevOps can then spend more time expanding their infrastructure, rather than catering for it. Many different areas have been enhanced. For example security, operations, fail-over, observability, failure detection, consistency, split-brain protection and primary election and flexible replication workflows.

This session highlights the new replication features in MySQL 8.0.

Those that were released pre and post-GA. Come and learn, directly from the engineers, how the new features help you operate, sustain and extend your MySQL Replication infrastructure.

Optimize and Troubleshoot MySQL using PMM 2

Peter Zaitsev, Percona

Optimizing MySQL performance and troubleshooting MySQL problems are two of the most critical and challenging tasks for MySQL DBAs. The databases powering your applications need to be able to handle changing traffic workloads while remaining responsive and stable so that you can deliver an excellent user experience. Further, DBAs are also expected to find cost-efficient means of solving these issues.

In this presentation, we will demonstrate how using recently released PMM version 2 enables you to solve these challenges, which is built on free and open-source software. We will look at specific, common MySQL problems and review.

MongoDB Analysis with Prometheus and Grafana

Akira Kurogane, Percona

Prometheus and Grafana are spreading through DevOps organizations around the world. We'll show how they been put together in PMM2 to create MongoDB monitoring timeseries graphs and alert triggers equal to any of the commercial solutions that currently exist.

- Quick demo of the classic DBA tasks
- Some new MongoDB DBA tricks
- The 'math' in MongoDB metrics:
 - » the base counters in the server process
 - » the prometheus timeseries metrics expressions
- Demonstration of fine-detail performance bottleneck analysis

How To Upgrade Like a Boss to MySQL 8.0?

Alkin Tezuysal, Percona

This talk will explain best practices for upgrade techniques in MySQL. In deep dive, we will go over how to upgrade successfully to MySQL 8.0. Explain MySQL 8.0 upgrade specific challenges. Go over gotchas and best practices. Review the latest version of MySQL 8.0 and bug reports.

Running Elasticsearch at Large Scale

Alex Cercel, Palantir Technologies

ElasticSearch is a high performance search engine based on Apache Lucene.

It is growing in popularity which means that more and more users are running part of their search patterns against the information stored on this datastore.

What is quite difficult to find in one place is a comprehensive set of best practices that will assist in efficiently running large scale clusters.

From my experience I've seen many bottlenecks created by some misconfigured parameter or rogue approach and I would like to offer some tips and tricks so others won't feel the same unnecessary pain. I've also seen how a tweek to a parameter nobody would have thought about, ended up to be a game changer.

In this 50 minutes talk I will cover tuning parameters and approaches that will assist the audience in running fast, efficient and well written queries on stable and performant clusters.

I will cover topics like schema design, talk about some different Garbage Collection mechanisms but also offer a good start on how to pick the right number of shards and replicas, role separation, what to look for when it comes to monitoring the cluster

and what to alert on to make sure that your cluster is stable and ... more.

All of this will be followed by some Q/A where I might be able to offer additional information that could help with other topics.

I look forward to welcoming as many people as possible to the wonderful world of search.

Managing MySQL at Scale in Facebook

Pradeep Nayak, Junyi Lu, Facebook

Section1: Terminology. Quick intro to shard/replicaset/instance/host and the scale of FB

Section2: Lifecycle of an instance

Section3: Tools that are used for moving instance between states: MPS Copy, rebuild_db etc

Section4: Shard movements(OLM) and touch base on balancing

Section5: Test Infrastructure: Shadow and Merlindb

Graph Databases: Introduction, Standardization, Opportunities

Peter Eisentraut, 2ndQuadrant

In this talk I want to present current developments in the world of graph databases and discuss how SQL databases could evolve to work well with graph data.

Graph databases are a major database paradigm next to relational databases, object databases, hierarchical databases, and so on. While there are popular standalone graph database products, some vendors from other fields including traditional SQL as well as NoSQL have begun to extend their products with graph data capabilities and interfaces. Relational databases have been successfully extended to support nonrelational uses such as document stores, so it is worth discussing whether and how graph database capabilities could also be integrated into SQL databases.

At the moment, the graph database world suffers from an overabundance of different query languages. There are, however, efforts underway to unify some of these and standardize them under the umbrella of standardization organizations such as W3C and ISO. This would make it easier for SQL database products, whose primary focus is not graph data, to add support for graph data query languages

without having to worry about competing languages and standards. I'll try to sort through this and show what is currently going on in graph database languages.

How to Instrument Your Code in Performance Schema

Georgi Kodinov, Oracle

Having all of the server code instrumented in performance schema helps users diagnose problems without having to use a debugger. But you don't have to be a server developer to instrument your code. MySQL 8.0 performance schema APIs allow good instrumentation of all code, including the code living in plugins and components. You can even be adding your own performance schema tables and expose your own volatile data and statistics into these. This session aims to walk you through the process of instrumenting your code that is to run inside the MySQL server process. And give you tips on how to do this most efficiently and where to go for further details.

Becoming Cloud Native: How Percona Brings Databases to Kubernetes Using Operators **Tyler Duzan, Percona**

This talk covers some of the challenges we sought to address by creating a Kubernetes Operator for Percona XtraDB Cluster and Percona Server for MongoDB, as well as a look into the current state of the Operators, a brief demonstration of their capabilities, and a preview of the roadmap for the remainder of the year.

This is a great way to find out more about the viability of running databases within a containerized environment, how you can self-service databases in a cloud-vendor agnostic way. You also get the chance to ask questions of the Product Manager and provide feedback on your experiences in the Kubernetes ecosystem.

The DB is Broken. Now What? **Carmen Mason, Vitalsource Technologies, Allan Mason, Ivan Groenewold, Pythian**

Sometimes the only piece of information that you are given is that the database is not working correctly. What do you do now? Panic? Call Mom? But, you're the DBA. What you do next is crucial.

In this talk, we'll cover the tools that you'll use to determine what other agents may be working against the database performance. We'll share our own experiences with crazy scenarios that stumped us, and the tools that we used to figure out the root cause.

Some examples include how systemd can ruin your day, NUMA imbalance, and issues unique to AWS, and GCP.

Automatic Upgrade and New Error Logging in MySQL 8.0 **Ståle Deraas, Oracle**

We will cover upgrade from MySQL 5.7 to MySQL 8.0 (8.0.15), going from legacy meta data storage to transactional data dictionary. We will cover the new possibilities for automation of upgrade, and the major advances upgrade speed and reliability as well as new consistency checks in the MySQL upgrade checker.

MySQL 8.0 has pluggable error log. We will talk about the traditional error logger and the JSON error logger, which empowers user with advanced filtering.

LIGHTNING TALKS

PostgreSQL Plan at execution time : A quick show

Jobin Augustine, Percona

A quick demo of decision makings happening during the execution phase in PostgreSQL latest versions.

JSON Array Indexes in MySQL

Norvald Ryeng, Oracle

MySQL 8.0.17 has the ability to create indexes over JSON arrays to speed up your queries. In this session we'll show examples and explain all you need to know about this new feature: How are array indexes created? How do they work? When are they used? Are there any limitations?

gdb basics for MySQL DBAs

Valerii Kravchuk, MariaDB Corporation

Sometimes MySQL DBAs have to work with stalled, hanged or unresponsive MySQL instance, where their usual SQL-based tricks do not work any more. Sometimes they can not even connect to check what's going on inside server.

In other cases they know what to do and everything still works, but they have to implement changes to read-only server variables. Server restart is often not an option in production, as it means some downtime and may cause negative performance impact.

In these cases one could do something given read and write access to server memory/internals. Here comes gdb, that, alone with careful reading of the source code helps to often resolve the problems described above. During this session I'll explain what can be done with gdb when server already is in troubles, and how to use gdb to "see" and understand MySQL internals better.

SPEAKERS

SPEAKERS

Ibrar Ahmed

Senior Software Engineer & PostgreSQL Consultant, Percona

Ibrar Ahmed is a Senior Software Engineer in Percona. Prior to coming to open source development, he had vast experience in software design and development.

The main focus was on the system level embedded development. After joining EnterpriseDB, an Enterprise PostgreSQL's company in 2006 he started his career in open source development specifically in PostgreSQL. He has contributed to the PostgreSQL community as well as other open source communities. His contribution is from the main performance feature enhancements to various PostgreSQL's modules. In the database field, he has experience in other well-known databases MySQL, Oracle, and NoSQL databases such as MongoDB and Hadoop. His experience is not limited to core databases, but with the tools related to databases like Hive, HBase and Spark. He also worked on integrating these tools with PostgreSQL. He joined Percona in the month of July 2018. He also authored multiple books on PostgreSQL.

[Join Heterogeneous Databases Using PostgreSQL Foreign Data Wrappers](#)

[Percona Distribution for PostgreSQL](#)

Manuel Arostegui

Database Administrator, Wikipedia

Manuel Aróstegui is a DBA working for the Wikimedia Foundation, the non-profit organization maintaining Wikipedia and other free knowledge collaborative projects.

[Backing up Wikipedia Databases](#)

Jobin Augustine

Senior Support Engineer, Percona

One of the people behind the PostgreSQL Distribution: BigSQL. Jobin has always been an active participant in the OpenSource communities, especially database communities and as an OpenSource advocate. He is one of the top bloggers in the PostgreSQL community. He is a contributor to various OpenSource Projects and he loves to code in C++ and Python in his free time. He worked for OpenSCG for 2 Years and Dell for 10 years as a Database expert.

[Why PostgreSQL is Becoming A Migration Target in Large Enterprises](#)

[PostgreSQL Plan at Execution Time: A Quick Show](#)

Sunny Bains

Director of Software Development, Oracle

Sunny Bains joined the Oracle/InnoDB team in 2006. He currently manages the InnoDB team. He has been writing software for longer than he can remember. Prior to joining InnoDB he wrote a database engine for a distributed database product at his previous employer. The networked database was used in a mission critical environment in the airline industry with very strict service level guarantees.

[InnoDB Management and Scalability Improvements in MySQL 8.0](#)

Charly Batista

Technical Lead, Percona

Charly worked as Java Architect for many years and using many different database technologies. He helped to design some of the features of the system used in the Brazilian Postal Service, the largest Java project in Latin America in that time. He also helped to design the database of the Brazilian REDESIM project, the system that is responsible for the municipalities taxation in Brazil. He now lives in China and works as Senior Engineer at Percona.

[10 Common Mistakes \(Java\)Developers Make When Writing SQL](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Krunal Bauskar

PXC Product Lead, Percona

Krunal has been working in MySQL space for over a decade now. He has actively contributed to MySQL community through different projects he has worked on.

He continues to contribute to Percona XtraDB Cluster making it one of the first choice when it comes to MySQL HA solutions. In past he was also associated with InnoDB team @ MySQL/Oracle where he authored some of the most important features of MySQL-5.7 like temporary table revamp, atomic truncate, undo log truncate, etc... He has also worked with Yahoo! Labs researching on bigdata problems and MySQL based database startup which is now part of Teradata.

[Percona XtraDB Cluster 8.0 \(PXC-8.0\)](#)

Satya Bodapati

Senior Software Engineer, Percona

Satya works with Percona Server Engineering team. He is responsible for Percona Server features, bug fixes, etc. He has an overall experience of 12 years in MySQL. He joined Percona in 2018. Before joining Percona, Satya worked with the InnoDB Development team at Oracle for 6 years. He has good knowledge of InnoDB internals and the new Data Dictionary. Loves to do C++ coding and write algorithms.

[Percona Server for MySQL 8.0: What It Is and How It Is Done](#)

Matthew Boehm

Senior Instructor, Percona

Matthew has been with Percona for almost 7 years. During that time he has fulfilled the roles of Senior Consultant, Senior Architect, and now as Senior Instructor, he delivers world-class training to clients all across the globe. His areas of speciality include Percona XtraDB Cluster, ProxySQL, MySQL Backups and Recovery.

[Getting Started with Kubernetes and Percona XtraDB Cluster](#)

John Boero

Channel Solutions Engineer, Hashicorp

John Boero is from Chicago living in London as a Channel Solution Engineer at Hashicorp. He helps educate Hashicorp's international partner network on Hashicorp technologies and enables them to sell to their customer base.

[Protecting your Secrets](#)

[Keynote](#)

Alex Cercel

Site Reliability Engineer, Palantir Technologies

Seasoned Linux, Windows, Network administrator, DBA, recently turned SRE with 15 years experience in the industry. For over 4 years I've been involved in running some of the largest implementations of MongoDB and Elasticsearch in the world. All round geek with a genuine passion for anything that runs on or involves computers. From the OS to games.

[Running Elasticsearch at Large Scale](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Colin Charles

Consultant, Galera Cluster

Colin Charles is a Consultant at Codership, the makers of Galera Cluster. Previously, Colin was on the founding team of MariaDB Server, and has been around the MySQL ecosystem including being an early employee at MySQL, and worked actively on the Fedora and OpenOffice.org projects. Colin has been a MySQL user since 2000.

[MariaDB Server 10.4: The Complete Tutorial](#)

Vicențiu Ciorbaru

Software Developer Team Lead,
MariaDB Foundation

Vicențiu Ciorbaru works for the MariaDB Foundation as a Server Developer and de facto team-lead of junior developers and new contributors. Vicențiu's past projects include Roles in MariaDB, Window Functions, Custom Aggregate functions. He also ensures MariaDB remains stable and up-to-date in Linux distributions by handling distribution specific bugs, especially in Debian.

[How a Modern Database Gets Your Data Fast: MariaDB Query Optimizer](#)

Michael Coburn

Product manager, Percona

Michael joined Percona as a Consultant in 2012, and progressed through various roles including Managing Consultant, Principal Architect, Technical Account Manager, and Technical Support Engineer, and is now leading the Product Management of Percona Monitoring and Management.

[Open Source Database Performance Optimization and Monitoring with PMM](#)

[What's New in PMM 2.0](#)

[Running PMM in Production at Tessi](#)

Jeremy Cole

Sr. Staff Production Engineer,
Shopify, Inc.

Jeremy Cole is a pioneer of MySQL scalability having worked for on MySQL for 19 years at MySQL AB, Yahoo!, Proven Scaling (providing advanced MySQL consulting), Twitter, Google, and now working as a Sr. Staff Production Engineer in the Datastores team at Shopify.

[MySQL on Google Cloud: The Good, The Bad, and The Ugly](#)

Stephane Combaudon

Senior DB engineer, Rakuten
Intelligence

Stephane has worked for almost 15 years in open source environments, and has spent most of the last 10 years solving performance or high-availability problems with MySQL and more recently MongoDB. He is now working for Rakuten Intelligence, a company that helps you keep track of your online purchases. His main role is to automate MySQL and MongoDB operations, and from time to time, he is also very excited to kill a nasty query or add a missing index on a 500GB table. When not in front of a computer, Stephane likes spending time with his family.

[Our journey to better MySQL Availability Using Global Transaction IDs, ProxySQL and Consul](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Jaime Crespo

Sr. Database Administrator,
Wikipedia

Jaime Crespo is a DBA working for the Wikimedia Foundation, the non-profit organization maintaining Wikipedia and other free knowledge collaborative projects. Lately his focus has been on recovery and database automation.

[Backing up Wikipedia Databases](#)

Tom De Cooman

SDM Team Manager, Percona

Tom has been active in Unix system's administration for 10+ years. Before joining Percona, he has been working for several consultancy companies and system's integrators, always working, focusing and offering solutions based on OSS.

[MySQL 101 Tutorial Parts 1 & 2](#)

Caleb James DeLisle

Protocol Designer, CJDNS

Caleb James DeLisle founded the cjdns Open Source project for cryptographically assured IP packet routing, he then went on to co-develop ncry.pt, the first Zero Knowledge pastebin using a secure key sharing technique which would later find it's way into ZeroBin, Mega.co.nz and others. During his time as a researcher at XWiki SAS, he authored the first realtime collaborative editor to make use of Nakamoto Blockchains for state synchronization and used this to launch the CryptPad project, a google docs like editor where the all content is encrypted and the server has no access to the keys. His Open Source work has been featured in Wired Magazine (USA), NewScientist (USA), Neue Zurcher Zeitung (Switzerland), Kommersant (Russia).

[Databases, Crypto & Decentralization](#)

Cristina DeLisle

Office and Legal Administrator, DPO
XWiki SAS

Cristina became interested in software when she began to work at XWiki and CryptPad, open source projects that made her more tech aware, coming from a legal background. She is a data protection officer and in this role she's following the GDPR's evolution, striving to ensure on a daily basis compliance with the applicable data protection rules.

[Data Protection and OSS in the Age of GDPR](#)

Ståle Deraas

Software Development Director,
Oracle

Ståle Deraas has been working with MySQL Database team since 2008. He is currently employed by Oracle, based in Norway. He is Senior Software Development Manager for the runtime team of the MySQL Database. He has a background in the database industry, working for Sun Microsystems on Java DB, and prior to that on Software Development Tools. He has a B.Sc Hons degree in Computing Science from the University of Glasgow.

[Automatic Upgrade and New Error Logging in MySQL 8.0](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Frédéric Descamps

MySQL Community Manager,
Oracle

“@lefred” has been consulting OpenSource and MySQL for over 15 years. After graduating in Management Information Technology, Frédéric Descamps started his career as a developer for an ERP under HP/UX. He will then opt for a career in the world of open-source by joining one of the first Belgian start-up dedicated 100% to free projects around GNU/Linux. It is in 2011 that lefred joined Percona, one of the leading MySQL-based specialists. He decided to join the MySQL Community Team in 2016 as a MySQL Community Manager for EMEA & APAC. Frédéric is also a regular speaker of OpenSource Conferences. His blog mostly dedicated to MySQL is lefred.be

[MySQL 8.0 InnoDB Cluster: Easiest Tutorial!](#)

[MySQL InnoDB Cluster: Advanced Configuration & Operation](#)

[MySQL Shell : The Best DBA tool? Extend the Shell with the New Extension Infrastructure](#)

Igor Donchovski

Lead Database Consultant, Pythian

MongoDB Certified DBA supporting the next-generation of database solutions in both MySQL and MongoDB. With a masters degree in Software Engineering, Igor enjoys the variety of challenges he faces while working at Pythian, rather than working on the same thing every day. When he's not working, he can be found biking

[Maintenance for MongoDB Replica Sets](#)

Tyler Duzan

Product Manager, Percona

Prior to joining Percona as a Product Manager, Tyler spent almost 13 years as an operations and security engineer in a variety of different industries. Deciding to take his analytical mindset and strategic focus into new territory, Tyler is applying his knowledge to solving business problems for Percona customers with inventive solutions combining technology and services.

[Becoming Cloud Native: How Percona Brings Databases to Kubernetes Using](#)

[Operators](#)

Yoav Eilat

Senior Product Manager, Amazon
Web Services

[Deep Dive on Amazon Aurora](#)

Peter Eisentraut

Deputy CTO, 2ndQuadrant

Peter has been a contributor to PostgreSQL development since 1999. He is a committer, core team member, and is involved in many aspects of the community. He works at 2ndQuadrant, a global service and support provider for PostgreSQL.

[Graph Databases: Introduction, Standardization, Opportunities](#)

Daniël van Eeden

Database Administrator,
Booking.com

Open Source enthusiast. Contributor to various projects in the MySQL ecosystem.
[Large Scale Deployment of SSL/TLS For MySQL](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Antonios Gianopoulos

Senior Database Administrator,
ObjectRocket by Rackspace

Antonios Giannopoulos is a Senior Database Administrator at ObjectRocket supporting MongoDB, Elastic and CockroachDB. Antonios has 18 years of experience in databases and system engineering. He really enjoys challenges in sharding and schema design and loves migrations from Relational to NoSQL world.

[New Indexing and Aggregation Pipeline Capabilities in MongoDB 4.2](#)

[What's New on Sharding in MongoDB 4.2](#)

Kathy Gibbs

Senior Database Specialist
Solutions Architect

[Accelerating Application Development with Amazon Aurora](#)

Pedro Gomes

Software Engineer, Oracle

Pedro is graduated and did his masters on Minho University where he did his dissertation on distributed systems. It was a mainly about the use of NoSQL databases in the industry and it included a ORM proof of concept. He joined Oracle on September 2012 after having worked for one year for the University Distributed System Lab doing some research.

[MySQL InnoDB Cluster: Advanced Configuration & Operation](#)

Jonathan Gourdon

Tessi

Jonathan started his career as a Software engineer in Bank IT. He joined Tessi in 2009.

As a constant learner, he progressed through various roles within the company and is currently working as an Automation Leader. His main role is to secure and automate solutions deployments, including monitoring, such as MySQL environments or NoSQL databases (Cassandra).

He works continuously towards the objective of providing production ready workflows.

In his free time, he loves climbing and being in nature.

[Running PMM in Production at Tessi](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Vinicius Grippa

Senior Support Engineer, Percona

Vinicius Grippa is a Percona Senior Support Engineer. Vinicius has a Bachelor's degree in Computer Science and has been working with databases for 10 years. He has experience in designing databases for mission-critical applications and in the last few years has become a specialist in MySQL and MongoDB ecosystems. Working in the Support team, he has helped Percona customers with hundreds of different cases featuring a very wide range of scenarios and complexities. Vinicius is also active in the OS community, participating in virtual rooms like Slack, and speaking at conferences, and MeetUps in Europe, North and South America.

[Open Source Database Performance Optimization and Monitoring with PMM](#)

[A Journey with MongoDB HA. From Standalone to Kubernetes Operator](#)

[Enhancing MySQL Security](#)

Ivan Groenewold

Lead Database Consultant, Pythian Group

Ivan has been supporting mission-critical environments for top of the line companies for over 15 years, while becoming a certified professional for Oracle, MySQL, MongoDB, Google Cloud and Amazon Web Services in the process.

Also co-developed the Database-as-a-Service platform for the biggest e-commerce site of Latin America (MELI:NSQ)

[The Database is Broken. Now What?](#)

Kenny Gryp

Principal Product Manager, Oracle

Kenny is a MySQL expert with a lot of experience in performance and architecture. While he has focused on MySQL and automation for over a decade, he has also developed expertise in other areas. He's a well-rounded, multidimensional nerd who is skilled in many things: making a coffee on the side of a mountain, electrical building codes for single family dwellings in the US, the perfect ratio of cumin and asafoetida in tadka dal, and providing occupational therapy to a shy rabbit named Winston. He's adept at talking, so don't feel the need to limit in-session questions to MySQL.

[MySQL 8.0 InnoDB Cluster: Easiest Tutorial!](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Robert Hodges

CEO, Altinity

Robert Hodges is CEO of Altinity, which offers software and services for ClickHouse, a high performance OSS data warehouse. His experience in data management includes relational database work at Sybase, developing SaaS applications on Oracle, and over 10 years working on replication and clustering products for MySQL. He was CEO of Continuent, Inc., when it was acquired by VMware in 2014. After 4 years working on multi-tenant cloud services, disaster recovery, and security he returned to data management at Altinity in 2019.

[ClickHouse Materialized Views: A Secret Weapon for High Performance Analytics](#)

Geir Høydalsvik

Software Development Director,
MySQL Server, Oracle

Geir Høydalsvik has been working with MySQL Database team since 2008. He is currently employed by Oracle, based in Norway. He is Senior Software Development Director and responsible for the development and maintenance of MySQL Database. He has a background in the database industry, working for the database startup company Clustra Inc. on the Clustra database and for Sun Microsystems on Java DB. He has a Master degree in Computer Science and a PhD in Software Engineering from the Norwegian University of Science and Technology.

[Keynote](#)

Guido Iaquinti

Operations Engineer, Slack

I'm a system engineer with academic background and experience in high volume/high availability Internet architectures. I'm a technology enthusiast excited about open source software. My passion is to develop, scale and automate complex systems.

[Strength in Numbers: Slack's Database Architecture](#)

Leigh Johnson

Staff Database Reliability Engineer,
Slack

Leigh Johnson is a Staff Database Reliability Engineer at Slack and a Google Developer Expert in Machine Learning. Her passion is automating herself out of a job, so she can spend all of her time with her corgi.

[Lessons From Building Automation for a Large Distributed Database](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Craig Kerstiens

Head of Product for Citus Cloud,
Microsoft

Craig currently works on product for open source databases at Microsoft. He joined Microsoft by way of Citus Data, where he ran product including their database as a service. Prior to joining Citus he ran product for Heroku Postgres. He's actively involved within the Postgres community, blogs about it regularly, and curates Postgres Weekly.

Keynote

Roman Khavronenko

COO, VictoriaMetrics

[VictoriaMetrics: Why and How We Built A Fast and Scalable Open Source Time Series Database](#)

Georgi Kodinov

Team Lead, MySQL Server General Team, Oracle

Georgi "Joro" Kodinov started working on MySQL in 2006 as a part of the optimizer team. He has moved to lead the Server General team, working on security, monitoring and the client-server protocol. Joro has background in banking IT.

[How to Instrument Your Code in Performance Schema](#)

Ameet Kotian

Staff Database Reliability Engineer,
Slack

Ameet Kotian is a Staff Database Reliability Engineer at Slack, where he is responsible for operating the infrastructure used to store all of Slack's data. Previously, he was one of the first site reliability engineers at Twitter, where he spent close to five years working on the graph storage service and Twitter's internal distributed database.

[Lessons From Building Automation for a Large Distributed Database](#)

Tibor Köröcz

Architect, Percona

Tibi joined Percona in 2015 as a Consultant. Before joining Percona, among many other things, he worked at the world's largest car hire booking service as a Senior Database Engineer. He enjoys trying and working with the latest technologies and applications which can help or work with MySQL together. In his spare time he likes to spend time with his friends, travel around the world and play ultimate frisbee.

[Percona XtraDB Cluster Tutorial](#)

Valerii Kravchuk

Principal Support Engineer,
MariaDB Corporation

Valerii Kravchuk supported MySQL users since 2005, in MySQL AB, Sun, Oracle, Percona and (since 2016) MariaDB Corporation.

[Tracing and Profiling MySQL](#)

[gdb Basics for MySQL DBAs](#)

Dimitri Kravtchuk

MySQL Performance Architect,
Oracle

Over 25 years in IT industry, mainly in Databases and Systems Performance. Since 2011 working full time on MySQL Performance at Oracle. Previous 15 years at Sun Benchmark Center in Paris. Previously at Informix Benchmark Team in France. Author of dim_STAT Performance Monitoring & Analyzing tool, actively publishing MySQL Performance related articles at <http://dimitrik.free.fr/blog>.

[MySQL 8.0 Performance: Scalability & Benchmarks](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Akira Kurogane

MongoDB Technical Lead, Percona

Akira Kurogane is an experienced database and search engine developer who advocates for long-term goals in the IT industry.

[What's New in Percona Server for MongoDB](#)

[Percona Backup for MongoDB](#)

[MongoDB Analysis with Prometheus and Grafana](#)

Pavel Lapaev

Product Manager, Tarantool

Background of finance and audit turned into several years as a technical trainer and telecom business trainer in Netcracker. Last two years works in Tarantool as a product manager.

[Tarantool: A NoSQL Database with SQL](#)

Junyi Lu

Production Engineer, Facebook

Junyi Lu is currently working at Facebook as a production engineer for the MySQL infrastructure. Responsible for developing automations which will manage tons of MySQL instances. Prior to this, he worked as a MySQL DBA at dianping.com, which is one of the biggest online-to-offline platforms in China.

[Managing MySQL at Scale in Facebook](#)

Aayushi Mangal

Support Engineer, Percona

Aayushi Mangal is a Support Engineer with Percona. Previously, she worked as a NoSQL Database Administrator for the retail based MNC. Her experience includes work as a database developer, database administrator, and as a web designer. Living in India, Aayushi is fond of nature, sightseeing and peace.

[Multi-Document Transactions Through MongoDB Shell](#)

Kevin Markwardt

Project Engineer, Pythian

Kevin Markwardt has twenty years of system administration experience ranging from MySQL, Linux, Windows, and VMware. Over the last six years he has been dedicated to MySQL and Linux administration with a focus on scripting, automation, HA, and cloud solutions. Kevin has lead and assisted with many projects focusing on larger scale implementations of technologies, including ProxySQL, Orchestrator, Pacemaker, GCP, AWS RDS, and MySQL.

Kevin Markwardt is a certified GCP Professional Cloud Architect, and a certified AWS Solutions Architect - Associate. Currently he is a Project Engineer at Pythian specializing in MySQL and large scale client projects. One of his new directives is Postgres and is currently supporting multiple internal production Postgres instances.

[Securing MySQL with Vault and Table Level Encryption](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Mykola Marzhan

Director of Server Engineering,
Percona

Mykola is Kubernetes and Clouds lover and, currently, his goal to bring databases into Kubernetes world. Since 2004, most of his career has focused on the development of monitoring, update and deployment systems.

[Comparison of Kubernetes Operators for MySQL](#)

Allan Mason

Database Consultant, Pythian
Group

Allan Mason enjoys constantly learning solutions to current and future issues; utilizing that information to help people and businesses. He works for Pythian doing exactly what he loves in a challenging environment. He is passionate about teaching and automation.

[The Database is Broken. Now What?](#)

Carmen Mason

Senior Database Administrator
VitalSource Technologie

Carmen Mason works at VitalSource Technologies as their Senior Database Administrator, where she helps decide on and plays with the latest cloud toys, and works with some pretty brilliant folks on a daily basis, and enjoys ROOT on her database instances. Her favorite thing about her job is the challenge and the people she works with.

[The Database is Broken. Now What?](#)

Yoshinori Matsunobu

Production Engineer, Facebook

Yoshinori Matsunobu is a Production Engineer and now a tech lead of MySQL team at Facebook. Yoshinori has been around MySQL community for over 10 years, and was a senior consultant at MySQL Inc since 2006 to 2010. Yoshinori created a couple of useful open source product/tools, including MHA (automated MySQL master failover tool), quickstack and MyRocks (RocksDB storage engine for MySQL).

[MyRocks and RocksDB Advanced Features and Performance](#)

Giuseppe Maxia

Software explorer, dbdeployer

Software Explorer working with a virtual team and maker of exploration tools.

Make difficult things easy.

MySQL community member for long, long time.

Four times receiver of MySQL community awards, twice as contributor, and twice as application of the year developer.

[Test Like a Boss: Deploy and Test Complex Topologies With a Single Command](#)

[Amazing Sandboxes with dbdeployer](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Boriss Mejias

PostgreSQL Consultant,
2ndQuadrant

I'm a holistic system software engineer, PostgreSQL consultant and trainer, free software activist, and headbanger. I have been working with PostgreSQL since version 9.1. First, as part of my job related to other projects, and now with full dedication since I joined 2ndQuadrant in 2017. I have a background as teaching-assistant and researcher on self-managing distributed systems, and I have presented in different conferences.

In open source, I have contributed to the Mozart-Oz programming language and Alfresco communities. In 2018 I have started the PostgreSQL User Group in Belgium.

[Pg_catalog Unrevealed! That Part of PostgreSQL You Are Probably Underusing](#)

Aleksei Milovidov

Team Lead, Yandex

Since 2008 I was working on development of data processing pipeline of Yandex Metrica - web analytics system.

Since 2015 I am responsible for development of ClickHouse - open-source column-oriented database management system, used for realtime analytical reports (<https://clickhouse.yandex/>).

I have 11 years experience with development of specialized data structures using C++ programming language.

[ClickHouse Features to Blow Your Mind](#)

Jim Mlodgenski

Principal Database Engineer,
Amazon Web Services

Principal database engineer at Amazon Web Services, Co-Chair - PostgresConf, Co-organizer of Philly PUG & NYC PUG, former CTO at OpenSCG

[Introduction to PL/pgSQL Development](#)

[Top 10 Mistakes When Migrating From Oracle to PostgreSQL](#)

Károly Nagy

Lead Database Engineer,
Salesforce

Karoly has been working with Salesforce since 2018 to improve the MySQL automation and help them scale the database infrastructure. Before that he was working with Dropbox and Booking.com on similar projects. His enthusiasm towards large scale distributed and autonomous systems had helped him developing time-tested frameworks to alleviate the day to day burden of operations.

[From Scheduled Downtime to Self-Healing in Less Than a Year](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Pradeep Nayak

Production Engineer, Facebook

Pradeep is a Production Engineer at Facebook and works with MySQL Infrastructure. He loves hacking code in python and builds bots to do things for him. While he is not working, he enjoys traveling and clicking pictures.

[Managing MySQL at Scale in Facebook](#)

John Nguyen

Senior Partner Solutions Architect

[Accelerating Application Development with Amazon Aurora](#)

Michal Nosek

Senior Solutions Engineer, Percona

During ten years of his career, Michal took different roles from a software engineer and business analyst to a technical sales consultant, always staying close to technology. He has hands-on experience with a broad range of programming languages and database technologies in different verticals.

As a Solutions Engineer at Percona, Michal's objective is to bridge the gap between business and engineering by providing customers with the best open-source software solutions and Percona services.

Michal holds a BS degree in Computer and Information Science from Cracow University of Technology. He lives in Gdansk, Poland with his wife. In his spare time, Michal enjoys traveling and sports.

[MySQL 101 Tutorial Parts 1 & 2](#)

Yuriy Olshanetsky

Senior DB Engineer, Rakuten Intelligence

Yurii started his career as DBA in Rakuten Intelligence 3 years ago. Rakuten Intelligence is a company that helps you keep track of your online purchases and save money. His main role is to automate MySQL and MongoDB operations, isolate DBs from the human factor, improve security and of course help people to make good queries. In free time Yurii prefers to be with his family and traveling.

[Our journey to better MySQL Availability Using Global Transaction IDs, ProxySQL and Consul](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Iwo Panowicz

Senior Support Engineer, Percona

Iwo is a Percona Support Engineer (EMEA) since August 2017. Prior to joining Percona Iwo was a System Administrator for more than 10 years with a focus on high availability, performance and automation. He also was helping many companies in designing infrastructure and ha/failover solutions.

[MySQL Has Gone Away: An In-Depth Look at the MySQL Networking Implementation](#)

Rachel Pescador

Product Marketing Manager,
Percona

Rachel Pescador is a Product Marketing Manager at Percona. She is responsible for highlighting Percona's open source database expertise by promoting Percona's support, consulting, managed services, and software solutions. Prior to joining Percona, Rachel spent 18 years in marketing and business development for some of the world's largest and most well-known professional services firms.

[What's New in PMM 2.0](#)

Vítor Pombeiro

Support Engineer, Microsoft

[MySQL, MariaDB & PostGreSQL as PaaS on Azure](#)

Alexander Rubin

Director, Data Architecture,
VirtualHealth

Alexander has over 10 years industry experience with the MySQL database and related technologies. His specialties are performance tuning, full text search, high availability, database infrastructure architecture and data warehouses. He has helped many MySQL customers design extremely high performance databases with optimized schema and queries.

[How to protect PII data in MySQL While Allowing Customers to Access the Database](#)

Norvald Ryeng

Software Development Senior
Manager, Oracle

Norvald has been at Oracle since 2011 and is the MySQL optimizer team lead, responsible for query optimization and execution. He holds a PhD in Computer and Information Science from the Norwegian University of Science and Technology.

[MySQL 8.0.1: Latest Updates](#)

[JSON Array Indexes in MySQL](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Art van Scheppingen

Database Engineer, Messagebird

Art van Scheppingen is Database Engineer at Messagebird. He's a pragmatic MySQL and Database expert with over 18 years experience in web development. He previously worked as a Senior Support Engineer at Severalnines. Prior to Severalnines he worked at Spil Games as Head of Database Engineering.

He has experience with a broad spectrum of (open source) databases: from MySQL to MongoDB, Cassandra to Hadoop and from Sphinx Search to SOLR. He regularly presents his work and projects at various conferences (Percona Live, Community Open House for MongoDB, FOSDEM) and related meetups.

[Benchmarking Should Never be Optional](#)

Sveta Smirnova

Principal Support Escalation Specialist, Percona

Sveta Smirnova is a MySQL Support Engineer with over than 10 years experience. Her main professional interests are problem-solving, working with tricky issues, bugs, finding patterns which can solve typical issues quicker, teaching others how to deal with MySQL issues, bugs and gotchas effectively. She is the author of the book "MySQL Troubleshooting" and JSON UDF functions for MySQL.

[Open Source Database Performance Optimization and Monitoring with PMM](#)

[Billion Goods in Few Categories: How Histograms Save a Life?](#)

Luis Soares

MySQL Replication Director, Oracle

Luis Soares is the MySQL Replication Director at Oracle. His research interests include replication technologies, dependable systems and high availability. Before joining the MySQL team, he was both a post-graduate student and a researcher at the University of Minho, Portugal, where he designed and implemented group-based replication protocols. Before that, he worked for a Portuguese multi-national company as a system analyst.

[MySQL 8.0: The New Replication Features](#)

Yura Sorokin

Senior Software Engineer, Percona

Yura is a Senior Software Engineer at Percona. He is a primary developer who implemented "Compressed Columns with Dictionaries" for Percona Server. Before joining in July 2015 he was leading backend dev team for a cloud file service with a focus on client-side encryption. He has 15+ years of software development experience, primarily in C++. Yura holds Master degree in Computer Science from National Technical University of Ukraine. He lives in Kyiv, Ukraine.

[Percona Server for MySQL 8.0: What It Is and How It Is Done](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Sugu Sougoumarane

CTO, PlanetScale

Sugu is the co-creator of Vitess, and has been working on it since 2010. Prior to Vitess, Sugu worked on scalability at YouTube and was also part of PayPal in the early days. His recent interest is in distributed systems and consensus algorithms. He occasionally shares his thoughts on his blog ssougou.blogspot.com.

[Vitess: Running Sharded MySQL on Kubernetes](#)

Feike Steenbergen

Software Engineer, Timescale

Feike has been working with PostgreSQL for over a decade. Performance related issues have been a recurring theme; whether working in retail (Zalando, bol.com), finance (Adyen) or now with TimescaleDB, whether with OLTP or with OLAP databases.

He's one of the organizing members of The Netherlands PostgreSQL Meetup.

Feike was part of the team that built Patroni, allowing fast failovers of PostgreSQL without human intervention. He has given PostgreSQL related trainings almost every year in the past 7 years.

[Parted Ways with Partitioning? It's Time to Reconsider](#)

Alkin Tezuysal

Sr. Technical Manager, Percona

Alkin Tezuysal is a Senior Technical Manager in Percona's Managed Services team. He has extensive experience in enterprise relational databases working in various sectors for large corporations. With more than 20 years of industry experience, he has acquired skills for managing large projects from the ground up to production. For the past 10 years, he's been focusing on e-commerce, SaaS and MySQL technologies. He managed and architected database topologies for high volume site at eBay Intl. He has several years of experience in 24X7 support and operational tasks as well as improving database systems for major companies. He has led MySQL global operations team on Tier 1/2/3 support for MySQL customers. In 2016 he joined Percona's expert technical management team.

[How to Upgrade Like a Boss to MySQL 8.0?](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Morgan Tocker

Community Development Manager,
PlanetScale

Morgan Tocker is the Community Development Manager at PlanetScale. He has worked in the MySQL space for the last 15 years, and was previously Product Manager for the MySQL Server at Oracle.

[My First 90 Days With Vitess](#)

Adamo Tonete

Senior Support Engineer, Percona

Adamo joined Percona in 2015, after working as a MongoDB/MySQL Database Administrator for three years. As the main database member of a startup, he was responsible for suggesting the best architecture and data flows for a worldwide company in a 7/24 environment. Before that, he worked as a Microsoft SQL Server DBA in a large e-commerce company, mainly on performance tuning and automation. Adamo has almost ten years of experience working as a DBA and in the past five years, he has moved to NoSQL technologies without giving up relational databases.

[A Journey with MongoDB HA. From Standalone to Kubernetes Operator](#)

[MongoDB Data Security: Custom Rules and Views](#)

[What's New in Percona Server for MongoDB](#)

Valentin Traën

Administrateur Bases de données,
Tessi

Valentin has been working for Tessi for more than 4 years, and as the DBA for 2 years. He is mainly responsible for the maintenance of Tessi's databases but the person who gets the job when a new query or table is needed. He is studying an engineering course at the famous French school, the CNAM and will graduate in 2020.

Valentin specializes in MySQL database engineering but also works with other database systems (Cassandra, PostgreSQL, Oracle, SQL Server ...).

I currently live in France, in Limoges.

[Running PMM in Production at Tessi](#)

Ben Tullis

Head of Infrastructure,
OpenCorporates Ltd

With over 20 years' experience focused on Linux systems management and Open Source infrastructure projects behind him, Ben has for the last six years played a key role in crafting the platforms underpinning OpenCorporates, the largest open database of companies in the world.

[OpenCorporates: Providing Transparency for the Public Benefit](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Marco Tusa

Consultant, Percona

Marco Tusa had his own international practice for the past twenty eight years. His experience and expertise are in a wide variety of information technology and information management fields, cover research, development, analysis, quality control, project management and team management. Marco is currently working at Percona as High Availability Practice Manager, previously working at Percona as manager of the Consulting Rapid Response Team on October 2013. He has been working as employee for the SUN Microsystems as MySQL Professional Service manager for South Europe., and previously in MySQL AB. He has worked with the Food and Agriculture Organization of the United Nation since 1994, leading the development of the Organization's hyper textual environment. Team leader for the FAO corporate database support. For several years he has led the development group in the WAICENT/Faoinfo team. He has assisted in

defining the Organization's guidelines for the dissemination of information from the technology and the management point of view. He has participated in field missions in order to perform analysis, reviews and evaluation of the status of local projects, providing local support and advice. He had collaborates with MIT Media Lab (Massachusetts Institute of Technology laboratory) and FAO as Sustainable Information Technology for developing countries Specialist in relation with the FAO's Special Program for Food Security for Senegal.

[Improving Enterprise HA and Disaster Recovery Solutions](#)

[Fortify Your MySQL Data Security in AWS Using ProxySQL Firewall](#)

Carlos Tutté

Support Engineer, Percona

Carlos joined Percona on February 2018, being part of AMER support team.

From Montevideo, Uruguay graduated as a Compute Engineer on 2012. Working on IT since 2010 taking different roles and technologies, first as a developer, and then moving to DBA and sysadmin.

[Open Source Database Performance Optimization and Monitoring with PMM](#)

[PCI/DSS Compliance with MySQL: 2019 Edition](#)

Aliaksandr Valialkin

CTO, VictoriaMetrics

Go developer, fasthttp author, VictoriaMetrics founder and core developer.

[VictoriaMetrics: Why and How We Built A Fast and Scalable Open Source Time Series Database](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Avinash Vallarapu

PostgreSQL Technical Lead,
Percona

Avinash Vallarapu joined Percona in the month of May 2018. Before joining Percona, Avi worked as a Database Architect at OpenSCG for 2 Years and as a DBA Lead at Dell for 10 Years in Database technologies such as PostgreSQL, Oracle, MySQL and MongoDB. He has given several talks and trainings on PostgreSQL. He has good experience in performing Architectural Health Checks, Performance Tuning and Migrations to PostgreSQL Environments.

[PostgreSQL For Oracle and MySQL DBAs and For Beginners](#)

[High Availability and Automatic Failover in PostgreSQL](#)

[Handling Transaction ID Wraparound in PostgreSQL](#)

Dimitri Vanoverbeke

Senior Solutions Engineer, Percona

At the age of 7, Dimitri received his first computer, since then he has felt addicted to anything with a digital pulse. Dimitri has been active in IT professionally since 2003 in which he took various roles from internal system engineering to consulting. Prior to joining Percona, Dimitri worked as a Open Source consultant for a leading Open Source software consulting firm in Belgium. During his career, Dimitri became familiar with a broad range of open source solutions and with the devops philosophy. Whenever he's not glued to his computer screen, he enjoys travelling, cultural activities, basketball and the great outdoors. Dimitri is living with his girlfriend in the beautiful city of Ghent, Belgium.

[MySQL 101 Tutorial Parts 1 & 2](#)

Josh Varner

Database Reliability Engineering
Manager, Slack

[Strength in Numbers: Slack's Database Architecture](#)

Markus Winand

SQL Renaissance Ambassador,
winand.at

Markus Winand is the SQL Renaissance Ambassador. He is on a mission to introduce developers to the evolution of SQL in the 21st century. His first book "SQL Performance Explained" is available in five languages and can be read for free on his website use-the-index-luke.com. Markus is currently working on his second book that is likewise available free of charge online at modern-sql.com. Markus can be hired as trainer, speaker and consultant via winand.at.

More resources: use-the-index-luke.com/

[More Than a Query Language: SQL in the 21st Century](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)

SPEAKERS

Matt Yonkovit

CXO, Percona

Before joining Percona, Matt worked at MySQL and Sun Microsystems as a Solution Architect, helping to build architectures for Fortune 500 and top Alexa-rated companies. Since the mid-'90s, he has worked with relational databases as a DBA, system administrator, and all around techie. During his career as a DBA, Matt has designed and supported massive database deployments that ran on MySQL, Oracle, DB2, SQL Server, Sybase, Postgres, and the AS/400.

Matt has a passion for performance and spends his free time doing things like benchmarking, finding outside-the-box solutions for today's performance challenges and trying to make things just go faster.

[Keynotes & MC](#)

Alexander Zaitsev

CTO, Altinity

Alexander Zaitsev is a co-founder and CTO of Altinity. He has been involved in software development, alongside academic research, since 1997. Alexander's interests include distributed architecture, databases and analytics. His focus is on building analytics solutions using database management systems capable of processing petabytes of data, such as Vertica and ClickHouse. Alexander has a Master's degree in mathematics and computer science from Lomonosov Moscow State University.

[ClickHouse for Time-Series Real-time Analytics](#)

Peter Zaitsev

CEO, Percona

Peter co-founded Percona in 2006, assuming the role of CEO. Percona helps companies of all sizes maximize their

success with MySQL. Percona was named to the Inc. 5000 in 2013. Peter was an early employee at MySQL AB, eventually leading the company's High-Performance Group. A serial entrepreneur, Peter co-founded his first startup while attending Moscow State University where he majored in Computer Science.

As CEO of Percona, Peter enjoys mixing business leadership with hands-on technical expertise. Peter is the co-author of High-Performance MySQL published by O'Reilly, one of the most popular books on MySQL performance. Peter blogs regularly on MySQLPerformanceBlog.com and speaks frequently at conferences.

Peter lives in North Carolina with his wife and two children. In his spare time, Peter enjoys travel and spending time outdoors.

[Keynote](#)

[InnoDB Architecture and Performance Optimization Tutorial for MySQL 8](#)

[Optimize and Troubleshoot MySQL Using PMM 2.0](#)

[SEE THE WEBSITE FOR ANY LATE CHANGES](#)



PERCONA
LIVE EUROPE
AMSTERDAM

30 Sept - 2 Oct 2019

CONFERENCE AND TUTORIAL **SCHEDULE**

