

**Postgres@CERN - PostgreSQL
Meetup at CERN - January
13th**

Report of Contributions

Contribution ID: 1

Type: **not specified**

Introduction

Friday, 13 January 2023 14:00 (25 minutes)

Presenter: DAFONTE PEREZ, Eva (CERN)

Contribution ID: 2

Type: **not specified**

Citus Data: The database distributed without complex

Friday, 13 January 2023 14:30 (50 minutes)

Interested in the issues of High Availability and Quality of Service, he contributes to the various projects he uses and promotes.

Citus Data is an extension of PostgreSQL that extends this famous relational database server to design distributed databases.

Distributed calculations, multi-tenant with isolated customers, resharding without downtime, ...

If this already sounds familiar, you'll learn more. And if not, you will finally know what a multi-tenant architecture is... And how easy sharding is!

Presenter: VILLEMMAIN, Cédric (Data Bene)

Contribution ID: 3

Type: **not specified**

Patroni + pgBackRest : Hero + Hero = Superhero?

Friday, 13 January 2023 15:50 (40 minutes)

The first hero gives you high availability and reliability. The second hero gives you disaster recovery and recovery from human errors. What happens if we combine those heroes to get a kind of superhero?

We're talking about Patroni and pgBackRest here: Can/will they integrate? We'll see that live, reproducible for you with easy step by step instructions.

Presenter: GUGEL, Julia (DBI Services)

Contribution ID: 4

Type: **not specified**

PostgreSQL: The Time-Series Database You Actually Want

Friday, 13 January 2023 16:50 (40 minutes)

Time-series data, or data being associated with its respective time of occurrence, is everywhere. From the obvious cases, such as metrics, observability, IoT data, all the way to logs, invoicing, or payment records.

While storing some of these in relational databases is standard practice, people often reach for specific time-series databases when volume gets high. But imagine if you could have all of them in the same database: PostgreSQL.

Join me for this session to learn more about the different types of time-series data and have a look at the naive, the native, and the scalable approaches to storing it in PostgreSQL. We'll contrast their usability and performance characteristics and show you why Postgres is the only database you need!

Primary author: ENGELBERT, Chris (Timescale)

Presenter: ENGELBERT, Chris (Timescale)

Contribution ID: 5

Type: **not specified**

YugabyteDB: Distributed SQL Open-Source PostgreSQL-compatible

Friday, 13 January 2023 17:35 (20 minutes)

A short introduction to Distributed SQL databases, to scale-out OLTP with cloud-native High Availability, Elasticity, and geo-distribution

Primary author: PACHOT, Franck (Yugabyte)

Presenter: PACHOT, Franck (Yugabyte)