

# Proceedings of the 2nd International Workshop on Composable Data Management Systems, co-located with 49th International Conference on Very Large Data Bases, 2023

Satyanarayana R Valluri, Mohamed Zait

*Databricks, Inc.*

## Workshop Chairs

- Satyanarayana R Valluri, Databricks Inc.
- Mohamed Zait, Databricks Inc.

## Steering Committee

- Wolfgang Lehner, TU Dresden, Germany
- Eugene Kogan, SingleStore
- Mark Callaghan, Meta Platforms Inc.
- Pedro Eugenio Rocha Pedreira, Meta Platforms Inc.

## Program Committee

- Shrikanth Shankar, Databricks, Inc.
- Renata Borovica, University of Melbourne, Australia
- Jacques Nadeau, Sundeck
- Jonathan Goldstein, Microsoft
- Stefan Krawczyk, DAG Works
- Walaa Eldin Moustafa, LinkedIn
- Sam Lightstone, Meta Platforms Inc.
- Garret Swart, Oracle
- Khaled Yagoub, Snowflake
- Philippe Bonnet, IT University of Copenhagen, Denmark
- Faisal Nawab, University of California Irvine

## Preface

Following a successful debut in 2022, the second edition of the International Workshop on Composable Data Management Systems (CDMS) will be held again this year along with the VLDB conference in Vancouver (Canada) on August 28th 2023.

The CDMS workshop seeks to bring together database researchers, practitioners, developers and users from academia, business and industry to discuss and explore the area of building database management systems with emphasis on composability and reusability.

While the format has changed to offer a more diverse content, the goals and themes of the workshop remain the same. For example, this year the schedule includes a very exciting panel discussion and 6 lightning talks along with 4 keynotes, 2 invited talks and 5 short papers.

This year's keynotes feature speakers with diverse backgrounds in the database industry: Orri Erling (Software Engineer, Meta Inc.), Mosha Pasumansky (CTO, Firebolt), Benjamin Wagner (Engineering Manager, Firebolt), Jordan Tigani (Co-Founder & CEO, MotherDuck), and Nikita Shamgunov (CEO, Neon Database). Orri Erling will share his experience in building a composable query engine in collaboration between Meta and academic partners and exploiting hardware to accelerate query execution. Mosha Pasumansky and Benjamin Wagner will present their experience reusing and replacing components of the DBMS stack and connect it with other efforts in the industry and academia. Jordan Tigani will discuss the hybrid execution system based on DuckDB that his team built at MotherDuck, but also discuss some further query topologies that are enabled by an execution model where a query is decomposed into parts some of which is executed locally and some executed remotely. Nikita Shamgunov will discuss his experience building an architecture where storage and compute are disaggregated, and use it to deliver Serverless Postgres in the cloud, and scale it up and down with the load without input or action from the platform user.

As stated earlier, this year's edition includes a very exciting panel titled: The Rise of Open Table Formats: Diving into the Next Decade of Data Lakes. The panel participants include three different camps: (i) Creators or representatives of the three most popular table formats: Rahul Potharaju (Delta), Ryan Blue (Iceberg), and Nishith Agarwal (Hudi), (ii) Builders of query engines that exploit the capabilities of those table formats: Dain Sundstrom (Trino) and Justin Levandoski (BigQuery) and (iii) Creators of benchmarks to evaluate the capabilities and performance of the table formats: Jesus Camacho Rodriguez (Microsoft GSL).

To ensure visibility of the workshop to people who did not attend, we will make the proceedings available through the proceedings of Joint Workshops at 49th International Conference on Very Large Data Bases (VLDB'23) and DBLP.

We would like to thank all the members of the CDMS Steering Committee and Program Committee for their support and hard work in putting together such a high-quality program. In closing, we welcome you to the CDMS 2023 workshop in Vancouver (Canada) and hope you will have a stimulating experience that will inspire you to contribute your own ideas next year or find something you can use back at your academic institutions or businesses.

*August 2023*

**Satyanarayana R Valluri**  
**Mohamed Zait**

# Contents

## Keynotes

- **Horizons of Composability**, *Orri Erling, Software Engineer, Meta Inc.*
- **Unbundling of the DBMS stack**, *Mosha Pasumansky, CTO, Firebolt and Benjamin Wagner, Engineering Manager, Firebolt*
- **Hybrid Query Execution; What is a database client, anyway?**, *Jordan Tigani, Co-Founder & CEO, MotherDuck*
- **Taking Postgres into the 21st Century**, *Nikita Shamgunov, CEO, Neon Database*

## Invited Talks

- **Innovating with storage formats to push the limits in a hybrid database**, *Eugene Kogan, Software Architect, SingleStore*
- **Holistic Extensibility for Integrated Data Analysis Pipelines in DAPHNE**, *Patrick Damme, Postdoctoral Researcher, Technische Universität Berlin*

## Short Papers

- **Relation-Based In-Database Stream Processing**, *Christian Winter (Technical University of Munich), Thomas Neumann (Technical University of Munich), Alfons Kemper (Technical University of Munich)*
- **The Gluten Open-Source Software project - modernizing Java-based query engines for the Lakehouse era**, *Akash Shankaran (Intel), George Gu (Intel), Weiting Chen (Intel), Binwei Yang (Intel), Chidamber Kulkarni (Intel), Mark Rambacher (Intel), Nesime Tatbul (Intel), David E. Cohen (Intel)*
- **Wisent: An In-Memory Serialization Format for Leafy Trees**, *Hubert Mohr-Daurat (Imperial College London), Holger Pirk (Imperial College London)*
- **Building a serverless Data Lakehouse from spare parts**, *Jacopo Tagliabue (NYU and Bauplan), Ciro Greco (Bauplan), Luca Bigon (Bauplan)*
- **E-Scan: Consuming Contextual Data with Model Plugins**, *Viktor Sanca (EPFL), Anastasia Ailamaki (EPFL & Google)*

## Lightning Talks

- **Polygraph - An Arrow Format for Large Apache Calcite Query Plans and Multi-Language Data Systems**, *Adam Kennedy, Voltron Data*
- **Composable Gateway for Databases**, *Aditya Jalan, Meta*
- **Using Multiple Composable, Hardware-Accelerate Executors**, *Felipe Aramburu, Voltron Data Inc.*
- **The Case for Decoupled Transaction Managers**, *Hiroyuki Yamada, Scalar, Inc.*
- **Technics in Accelerating Query Processing on GPU**, *Jimmy Lu, Meta*
- **Towards Efficient and Secure UDF Execution with BabelFishLib**, *Philipp Grulich, Technische Universität Berlin*

## Panel Discussion

- **The Rise of Open Table Formats: Diving into the Next Decade of Data Lakes**, *Walaa Eldin Moustafa, Senior Staff Software Engineer, LinkedIn*