

Hall 1

Hall 2

Hall 3

How to Teach LLM to Work with Data Instead of Just Writing Plausible SQL
Maksim Statsenko Yandex

11:15-12:00

What Happens Between SELECT and DATA
Petr Gurinov Yandex Cloud

11:15-12:00

How Platformization and AI Are Changing the Analytics Development Lifecycle: T-Bank's Experience
Dmitrii Rudnev T-Bank

11:15-12:00

LLM Under Load: How to Measure the Performance of Self-Hosted Models
Roman Peskov Cian

12:30-13:15

Kafka News: KRaft, Queue, Tiered Storage (And a Bit About YDB)
Andrey Serebryanskiy Yandex Cloud

12:30-13:15

S3, HDFS, POSIX... or All of the Above? Building a Data Lake the Chinese Way with CubeFS
Ivan Arkhipov

12:30-13:15

State of Iceberg REST Catalogs: What We're Missing and How to Make a DIY Control Plane
Vitaliy Moiseev Ostrovok!

14:10-14:55

Is There Life After dbt?
Alexandra Popova Positive Technologies

14:10-14:55

PostgreSQL Performance Diagnostics, or the Detective Called "Something's Slowing Down the Database"
Stepan Fomichev Yandex Cloud

14:10-14:55

LLM Ops: Optimization of Inference and ML-serving in a Real Production Cluster
Dmitry Ibragimov Lemana Tech

15:25-16:10

Data Marts on Data Lakehouse: A Major Migration from Greenplum 6
Artemii Naumov Lemana Tech

15:25-16:10

Knowledge Graph as an Infrastructure for AI Agents: From Datasets to a Single Graph
Aleksandr Nepochatykh Sber

16:40-17:25

Vector Search in PostgreSQL: pgvector Under the Hood
Daria Barsukova Postgres Pro

16:40-17:25

Datapipe — Data Transformation Using K8s and S3
Sergey Zakharchenko EPOCH8

16:40-17:25

Buying More or Keep On Using: How an AI Assistant Became Ready for Production
Alexander Podmoskovniy Flant

17:55-18:40

Migration of Data Management Tools to OMD at Magnit
Oleg Molchanov Magnit

17:55-18:40

18:40 - 22:00

Networking

Hall 1

Hall 2

Hall 3

Postmortem Comparisons of Agentic and Classical AutoML: Typical Pitfalls of the Agentic Approach
Valeriia Dymbitskaia Uppini

11:00-11:45

Data Streaming Lakehouse: How to Stream Data Into Paimon and Not Drown
Kirill Romanikhin Place.01

11:00-11:45

The MDM That Stores Nothing: How to Match Data Without Centralizing It
Iurii Goryntsev Arenadata Catalog

11:00-11:45

Reading Faster Than Ceph Can Serve: How We Built S3 Sharding With No Extra Infrastructure
Dmitrii Listvin Avito

12:15-13:00

From Text-to-SQL to Trusted Analytics: Building an On-Prem Semantic Layer for AI Agents
Igor Dmitriev Independent expert

12:15-13:00

Transactions in PostgreSQL: Parallelizing Non-Parallelizable
Daniil Davydov Postgres Professional

12:15-13:00

YTsaurus in the Wild: Pros, Cons, and Pitfalls
Nikita Blagodarnyi Chestny znak

13:30-14:15

Sketches: Useful in Practice or Just Amazing Mathematics?
René van Bevern

13:30-14:15

Why the Future of AI Is "Vectorless" and How We Tested It with the Operator Assistant
Andrey Nosov Raft

13:30-14:15

Pimp My Ride: Adapt Your Old Retriever to New Challenges
Vladislav Popov Tochka Bank

15:15-16:00

How to Search for a Memory Leak in the Storage for a Month and Find Out That It Actually Does Not Exist
Mikhail Motylenok YADRO

15:15-16:00

NiFi Threads Review and Deploy Via Git
Klavdia Popova Sibur Digital

15:15-16:00

ML Against Hackers. Processing Hundreds of Thousands of Events Per Second
Nikolai Lyfenko Positive Technologies

16:30-17:15

Data Contracts: When a Schema Becomes a Contract
Nikita Borzunov Uzum Market

16:30-17:15

Metric Store as a Boost for AI
Dmitriy Shirokov Yandex Taxi

16:30-17:15