Eduard Vlasov

Staff Software Engineer | Data Platforms and Distributed Systems

Amsterdam, Netherlands

driver.edro@gmail.com | linkedin.com/in/justedro | github.com/justedro | eduard.vlasov.nl

Summary

Staff Software Engineer focused on distributed systems, infrastructure, and large-scale data platforms. Combines hands-on systems engineering with cross-team technical leadership across architecture, platform design, and production optimization. Experienced in turning complex technical challenges into reliable, efficient, and scalable systems, including infrastructure for AI-enabled applications.

Experience

Miro Amsterdam, Netherlands

Staff Software Engineer *March 2022 -- Present*

- Designed and rolled out a data-collection standard across 15+ product domains, reducing integration effort from two weeks to a single commit; documented for engineers and coding agents
- Co-designed Miro's company-wide Pub/Sub event bus on Kafka and an in-house SDK
- Evaluated multiple self-hosted LLMs and built a cost-estimation tool comparing self-hosted inference costs with hosted AI provider pricing to inform infrastructure decisions
- Building an internal knowledge aggregation system that unifies information from multiple company sources and enables AI agents to run structured queries over a knowledge graph
- Owned Miro's event-tracking platform, processing 5B events per day and used by most teams at Miro; reduced analytics latency from daily batches to approximately 2 hours
- Technical lead for an 11-engineer data team, collaborating with about 30 engineers across adjacent teams
- Expanded technical ownership across multiple infrastructure domains; began reporting directly to the Head of Infrastructure in July 2026

Senior Data Engineer *December 2019 -- March 2022*

- Scaled event collection through 10x traffic growth in 2020 and designed analytics ingestion across sharded PostgreSQL, Redis, and DynamoDB during the monolith split
- Implemented streaming and batch pipelines on Kafka, Spark (Databricks), Airflow, Snowflake, and dbt

Data Engineer *March 2018 -- December 2019*

- Built the clickstream collection system and data pipelines; managed data infrastructure, data models, BI, and analytics

ER-Telecom Holding Russia

Oracle Database Developer *March 2012 -- March 2018*

- Developed core billing modules in Oracle for a major Russian broadband operator with millions of subscribers
- Tuned complex PL/SQL procedures and SQL queries over very large data volumes in critical billing workloads
- Introduced new technologies and development practices to the team and built internal developer tooling

Skills

Programming: Java, Kotlin, Scala, Python, SQL, PL/SQL, Bash, C, C++, Rust
Data & AI: Kafka, Spark, CDC, Pub/Sub, LLM evaluation, self-hosted inference, cost modeling, knowledge graphs
Databases: PostgreSQL, Redis, DynamoDB, Oracle Database, database sharding
Cloud and Infrastructure: AWS, Kubernetes, Terraform, Docker, Helm, Ansible, Linux, Git, CI/CD
Engineering: Distributed systems, system design, performance tuning, operating systems, computer architecture
Languages: English (full professional), Russian (native)

Awards

World IT-Planet 2016/17, 1st place in SQL, June 2017 | ACM ICPC, regional semi-finalist several times, 2010 -- 2013

Education

Perm National Research Polytechnic University *Perm, Russia*

Specialist degree (2009--2014) and Master's degree (2014--2016), Computer Software Engineering