

Sergey Polovko

Principal Software Engineer · Storage, Databases & Distributed Systems

Vancouver, BC, Canada · sergey@polovko.me · polovko.me · github.com/jamel · linkedin.com/in/polovko

SUMMARY

Principal software engineer with nearly 20 years of experience building high-performance storage, database, and observability systems — from Linux kernel block I/O to cloud-scale data warehouses. Currently at NVIDIA, working on data-path algorithms for NVMesh block storage. Previously built storage and metadata engines for Amazon Aurora and Redshift, and led the team behind Yandex Monitoring, ingesting 500M metrics per second. Starts org-wide technical initiatives and drives them to completion.

EXPERIENCE



NVIDIA · Remote (Vancouver, Canada)

Principal Software Engineer, NVMesh Block

Feb 2026 – Present

- Develop data-path algorithms — JBOD, mirroring, erasure coding — for [NVMesh](#), NVIDIA's software-defined NVMe block storage.
- Work across a Linux kernel block-layer module and user-space services, in C.



Amazon Web Services · Vancouver, Canada

Software Development Engineer III, Amazon Aurora & Amazon Redshift

Sep 2022 – Jan 2026

- **Aurora next-generation storage.** Designed and built from scratch a page-level storage engine with MVCC transactional semantics, integrated with Aurora's distributed storage nodes, in close collaboration with the PostgreSQL engine team.
- **Redshift zero-ETL with Aurora PostgreSQL.** Built the Aurora PostgreSQL-specific functionality on the Redshift side of the [zero-ETL integration](#) — generally available October 2024, replicating data into Redshift within seconds of it being written to Aurora.
- **Build and dependency modernization (self-initiated).** Started and drove the org-wide migration of Redshift builds and third-party dependencies to Amazon's standard build infrastructure. The audit surfaced vulnerable third-party dependencies; organized about 30 engineers to upgrade, replace or remove them, and reviewed and shepherded every change through release. Led a parallel Python 2.7 → 3 migration track with another 10–15 engineers. The team completed the migration the following year.

Software Development Engineer II, Amazon Redshift

Mar 2022 – Sep 2022

- **Metadata storage redesign.** Replaced the flat-array metadata layout — which required loading the entire superblock from Amazon S3 into memory — with a durable B⁺-tree that loads metadata on demand. Made cold queries faster and cheaper and let smaller clusters work with larger datasets.

- Led a team of up to 15 engineers building and operating Yandex's company-wide time-series monitoring platform (C++, Java): about 500 million metrics/s ingested, 150K queries/s.
- Designed a custom binary metrics format that cut network traffic 40× and server CPU and memory use 10×.
- Drove the launch of the internal system as [Yandex Cloud Monitoring](#), a public service for external customers that also monitors Yandex Cloud's own managed and infrastructure services.
- Designed an extensible monitoring agent: a C++ core with a plugin system for custom Python collectors.
- Led rewrites of the alerting and data-collection subsystems, and automatic shard rebalancing that moves load from overloaded hosts to hosts with spare capacity.
- Shipped a self-service admin UI so internal teams configure their own monitoring instead of filing requests; the team also ran Yandex's company-wide Grafana and Graphite installations.

- Built core components of YQL, the SQL dialect for Yandex's MapReduce platform YT: the C++ query-execution virtual machine and distributed, fault-tolerant query scheduling in Java. Used daily by thousands of engineers; YQL is now open source as the query language of [YDB](#) and [YTsaurus](#).

- Built a C++ analytical road-traffic service delivering real-time congestion data to media partners, including radio stations across Russia.

Earlier experience

- Lead Java Developer, Private Trade (Moscow, 2010–2012): scaled the platform with Terracotta, Nginx and MySQL tuning; built integration with the corporate data warehouse.
- Software Engineer at fintech startups (Novosibirsk, 2007–2010), alongside university studies.

OPEN SOURCE & TALKS

- Original author of the **YDB Java SDK** — github.com/ydb-platform/ydb-java-sdk
- *Architecture of a High-Performance, Highly Available Monitoring System at Yandex* — HighLoad++ Siberia 2019, Novosibirsk (in Russian): [abstract](#), [video](#)

SKILLS

Languages: C, C++, Rust, Go, Python, Java

Systems: Linux kernel (block I/O layer), storage engines, MVCC and transactions, PostgreSQL internals, distributed systems, time-series databases, networking

Performance: perf, eBPF, profiling and bottleneck analysis, concurrent and parallel programming

EDUCATION

 **Novosibirsk State Technical University** · Novosibirsk, Russia

B.Sc. in Computer Science and Engineering

2004 – 2010