

Stefan Vercillo

✉ svercillo7@gmail.com 📞 (917)-488-2282 🌐 stefanvercillo.com 🔄 svercillo 📄 svercillo

🎓 Education

University of Waterloo, BAsC Computer Engineering

Sep 2018 – May 2023

📁 Professional Experience

Palantir, Senior Software Engineer / Intern – Streaming 📄 Sep 2022 – Present | New York, NY

- Built a **B+ Tree** state module on top of **Flink**, enabling large-scale streaming interval and KNN joins with sorted output, secondary indexes, and memory-efficient eviction (**\$7M->\$1K** annual infra cost reduction in first production join instance)
- Designed a parallelized **Kafka** archival module that boosted throughput **75x** and enabled recovery from multi-terabyte backlogs within minutes, preventing critical disk outages and data loss across all production streaming pipelines.
- Achieved up to **100x faster** stream backlog processing by building and orchestrating a **Spark** job that reorganized multi-terabyte Kafka archives into a virtual bucket layout, eliminating partition skew and I/O bottlenecks.
- Built a custom parallel **Flink** ingest from virtual buckets, ensuring per-key ordering, fairness, and dynamic load balancing; **reduced replay times from months to hours** while maintaining consistency and fault tolerance guarantees.
- Developed an async request batching layer to reliably support high-latency, expensive operations at scale, enabling **real-time LLM**-based intelligence enrichment and AI transformations on high-volume video, audio, and imagery streams.
- Innovated streaming ingest method utilizing **auto-scaling** containers for flexible data population, enabling edge use cases such as live audio calls that seamlessly integrate with a live AIP-driven recommendation engine 📄
- Implemented metrics-driven dynamic scaling, ensuring **Flink** workloads maintain performance while minimizing unnecessary compute costs.

Twitter, Software Engineer Intern – Kernel and Encryption Team May 2022 – Aug 2022 | San Francisco, CA

- Deployed **Booster** 📄, a high-performance init-ram-fs generator across Twitter production Linux machines, reducing early-stage boot time by 15% while supporting full-disk encryption with TPM/network-bound unlock.

Federato AI, Early Software Engineer Dec 2021 – May 2022 | San Francisco, CA

- Engineered a secure multi-tenant secrets platform using HashiCorp Vault with webhook sidecars and per-client VPC isolation, automating VPN provisioning and access controls to centralize governance and strengthen compliance at scale.
- Designed and centralized continuous deployment workflows with GitOps-driven ArgoCD, enabling scalable, automated rollouts across client clusters and reducing operational overhead while ensuring consistent, auditable deployments.

TradeLogiq / Omega ATS, Quantitative Developer Jan 2021 – May 2021 | Toronto, ON

- Engineered a high-performance ingestion layer for TSX exchange data, ensuring consistent, real-time market data delivery.
- Optimized **UDP** retransmission in **C++**, cutting latency by 70% and improving order book accuracy under packet loss.

Enso Connect, Software Engineer Dec 2020 – May 2021 | Toronto, ON

TD Bank, Quantitative Developer Intern May 2020 – Sep 2020 | Toronto, ON

Sky View Suites, Software Engineer Sep 2019 – Jan 2020 | Toronto, ON

Ford Motor Company, Software Engineer Intern Jan 2019 – May 2019 | Waterloo, ON

📁 Projects

Database Systems Implementation, C++

- Created a relational query cost analyzer by evaluating materialization costs of temp tables, hash/tree index vs table scan selection costs, and considering byte reduction through projections and selection order using provided schema & statistics
- Developed a query tree visualizer and optimizer by intelligently exploring a subset of possible new indexes, rebalancing the query tree to prioritize pipelined joins, and maximizing byte reduction by strategically pushing selections and projections

🔧 Skills & Aptitudes

Libraries/Frameworks: Apache Flink, Hadoop, Apache Spark, Spring, Django, Cuda, React, Vue.js, UDP/TCP networking

Tools: Kafka, Kubernetes, Gradle, Maven, Docker, Terraform, ArgoCD, Git CI/CD, Redis, Kinesis, AWS, GCP, Azure

Languages: Java, Python, C++, Rust, SQL, GoLang