

MICHAEL MABINUOLA

Database/AI Engineer | Performance, Vector Search and Data Pipelines

Seoul, South Korea | michaelmabinuola@gmail.com | [LinkedIn](#) | [Portfolio](#)

SKILLS

Programming Python, SQL, PL/Python, PL/SQL, Bash, R, REST API Development, Rust (Basic)	Systems Linux, Docker, High availability clustering, Distributed systems, S3 Object Storage	Data Pipelines Apache Kafka, Airflow, dbt, ETL design, data validation	AI & Search Vector Databases, Semantic Search, RAG Pipelines, AI Agents, TensorFlow, machine learning
Databases SingleStore, MySQL, ClickHouse, MongoDB, OLAP/OLTP architecture	Linux Performance CPU, memory, page cache, filesystem	Visualization Tableau, Streamlit, Plotly	

EXPERIENCE

Database Engineer | 에이플랫폼 (Agile Platform) | Seoul, South Korea | Dec 2023 - Present

IT service and consulting company specializing in SingleStore DB.

[Responsibilities]

- Perform online and offline database migrations.
- Diagnose and resolve SingleStore performance bottlenecks for on-premise customers using query profiles, execution plans and OS-level metrics.
- Tune SingleStore for vector data storage and retrieval. Benchmark and evaluate vector index performance
- Build semantic search systems and RAG based AI pipelines, and provide technical support for AI search and vector functions
- Design Apache Kafka data pipelines that simulate customer workloads for proofs of concept (PoC)
- Write technical reports and documentation in English and including the company's reference guide on AI workloads in SingleStore.
- Evaluate competing database systems against SingleStore on architecture, performance, and features, and write comparison reports that support customer and internal decisions.

[Selected Projects]

- Unlimited storage PoC on VAST object storage - Validated VAST as S3-compatible storage for SingleStore: lifecycle tests (backup, HA, PITR, scaling), TPC-H benchmarks on hot and cold tiers, and throughput under load.
- Vector search with Matryoshka embeddings (2026) - Ran the benchmarks for a co-authored study published on the SingleStore blog: two-stage search on 10M vectors cut index memory 87% and raised throughput 6.6x at 96% recall ([Vector Search with MRL](#))
- Query profile analysis tool (in progress) - Building an internal tool that analyzes SingleStore query profiles in on-premise environments, for use by the support team.
- Hamming distance for SingleStore (2026) - Built a Rust WebAssembly UDF that adds Hamming distance over binary vectors, which SingleStore lacks natively, and published the guide on the company wiki ([SingleStore Hamming Distance](#))

Data Engineer / Data Analyst | 별따러가자 (Starpickers) | Seoul, South Korea | Sept 2021 – Nov 2023

Telematics company that fuses Inertial Measurement Unit (IMU) and GPS data to detect driving patterns.

[Responsibilities]

- Designed and implemented end-to-end ETL pipelines processing 1M+ rows daily from 300+ users.
- Built an SVM model that classified road vs. sidewalk driving with 92% accuracy.
- Developed algorithms to detect driving maneuvers: U-turns, rapid turns, overtaking

- Co-developed a driving behavior scoring system for usage-based insurance.
- Configured real-time workflow monitoring and alerting in Apache Airflow.
- Ran data quality validation to ensure accuracy and integrity, and maintained KPI monitoring dashboards.

EDUCATION

- Google Career Certificate in Advanced Data Analytics | Jan 2023 ~ July 2023
- MBA - Chung-Ang University (4.16/4.5) | Sept 2019 ~ August 2021
- BSc (Honors) Economics - Don State Technical University (4.83/5.0) | Sept 2014 ~ July 2018

CERTIFICATES/AWARDS

- SingleStore Certified Developer - 2024
- Junction Asia Certificate - 2022
- Data Analytics Certificate - 2020
- Government of Korea Scholar (GKS) - 2018
- Nigeria/Russia Bilateral Educational Scholarship - 2013

SPOKEN LANGUAGES

- English (Native), Russian (Proficient), Korean (Intermediate)