

KR

Kimberly Robinson

Data Engineer

Staff data engineer with thirteen years building data platforms for media and e-commerce. Most recent work has been on lakehouse architecture, governance, and the unglamorous job of getting 90 analysts to agree on definitions.

CONTACT INFORMATION



(718) 555-0119



david.okonkwo@example.com



Brooklyn, NY 12345

EDUCATION

- M.S. Computer Science, University of Illinois Urbana-Champaign, 2011
- B.S. Computer Engineering, University of Minnesota, 2009

KEY SKILLS

- Lakehouse architecture (Databricks, Delta Lake, Unity Catalog)
- Distributed processing with Spark, Trino, and Flink
- Streaming systems: Kafka, Kinesis, Pulsar
- Data modeling and dimensional design for analytics
- Data governance, lineage, and classification frameworks
- Pipeline orchestration with Airflow and dbt
- Cloud data platforms on AWS and Azure
- SQL performance tuning across Snowflake, Redshift, and Vertica
- Python and Scala for production data systems
- Cost optimization and capacity planning for petabyte-scale workloads
- Technical leadership, RFC authoring, and architecture review

PROFESSIONAL EXPERIENCE

Staff Data Engineer | Halberd Media Group, Brooklyn, NY 2021 - Present

- Tech lead for the lakehouse migration: Databricks + Unity Catalog replacing a Hadoop cluster and three regional warehouses.
- Drafted the company's data classification policy alongside legal; now governs ~2.1 PB across 14 business units.
- Cut critical-path pipeline failures from about 9 per week to under 2 by introducing tiered SLAs and a paging policy that distinguishes ad revenue jobs from internal reports.
- Chair the data architecture review board; review 6-10 design docs per month.
- Wrote the internal RFC that killed a half-finished Kafka-to-Elasticsearch project burning roughly \$1.6M annually in infra and headcount.

Principal Data Engineer | Sundial Commerce Co., Minneapolis, MN 2017 - 2021

- Architected the order events platform on Kinesis and Delta Lake; processed Black Friday peak of 1.4B events without a paged incident.
- Founded the data platform guild, a cross-team forum of 22 engineers that owned shared tooling decisions.
- Reduced warehouse query costs by ~38% through a workload migration to Trino for ad-hoc analytics.
- Hired and onboarded 7 data engineers across two teams; designed the interview loop still in use.
- Represented data engineering on the security review for the company's GDPR program.

Senior Data Engineer | Sundial Commerce Co., Minneapolis, MN 2014 - 2017

- Owned the customer behavior pipeline used by personalization and email marketing.
- Migrated a Hive-based ETL to Spark on EMR, cutting end-to-end runtime by half.
- Built the first internal data catalog (Python + Postgres + a small UI) before commercial options were viable.
- Mentored four engineers, two of whom now lead teams.

Data Engineer | Bramblewood Digital, Chicago, IL 2011 - 2014

- Built ETL for ad impression and click logs across three ad networks into a Vertica warehouse.

- Cross-functional partnership with legal, security, and analytics teams

- Wrote a deduplication library used in eight downstream pipelines.
- On-call lead for nightly batch; on-call documentation I wrote was used long after I left.
- Helped run the company's first data engineering interview track.