JULIAN MERCER

SENIOR DATA ENGINEER

San Francisco, CA | (415) 555-0148 | julian.mercer@email.com | linkedin.com/in/julianmercer

PROFESSIONAL SUMMARY

Senior Data Engineer with 11+ years designing lakehouse platforms, event pipelines, and governance for marketplace and social products. Leads migrations across teams, sets data standards, and mentors engineers while keeping cost and reliability in check.

CORE SKILLS

Apache Spark | Apache Flink | Kafka | Iceberg | Airflow | Trino | Python | Scala | SQL | Data contracts | Data governance | Lineage | Cost optimization | Technical leadership | Mentoring

PROFESSIONAL EXPERIENCE

Senior Data Engineer Feb 2021 - Present
Airbnb, San Francisco, CA

- Led a six-engineer migration of 400 Hive tables to an Iceberg lakehouse, finishing in nine months with no missed downstream deadlines.
- Designed the shared event-ingestion layer on Kafka and Flink that now serves 25 product teams at a 99.9% on-time delivery rate.
- Set data contract and ownership standards adopted by 12 teams, cutting breaking schema changes reaching analysts by more than half.
- Reduced warehouse and storage spend by 19% through tiering, compaction, and retiring unused datasets after a usage audit with finance.
- Mentor four engineers and run the quarterly design review for the data platform group, writing the template other teams copy.

Data Engineer Apr 2017 - Jan 2021
Pinterest, San Francisco, CA

- Rebuilt the ad-reporting pipeline on Spark and Airflow, bringing daily data availability forward by three hours for sales teams.
- Introduced anomaly alerts and lineage tracking across 60 critical jobs, which shortened incident diagnosis from hours to about 30 minutes.
- Partnered with machine-learning engineers to build a feature pipeline serving home-feed models, with backfills that stayed reproducible.
- Interviewed candidates and onboarded six new hires with a written first-month plan and pairing schedule.
- Cut a feature-store backfill from three days to eight hours by repartitioning Spark jobs and caching shared lookups.

Software Engineer, Data Jul 2015 - Mar 2017
Yahoo, Sunnyvale, CA

- Maintained Hadoop and Oozie workflows for advertising logs, processing several terabytes a day for reporting and billing teams.
- Moved a legacy Pig job to Spark, which shortened its nightly run by roughly 60% and cut cluster usage during peak hours.
- Wrote the operational runbook and alerting for the team's batch jobs after a data-loss incident, and ran the follow-up review.
- Built a log-validation job that compared source and loaded record counts hourly, and billing errors from missing data stopped in the next quarter.

EDUCATION

B.S. in Computer Engineering, University of Texas at Austin, Austin, TX

CERTIFICATIONS

Confluent Certified Developer for Apache Kafka | Google Cloud Professional Data Engineer