JULIAN MERCER

DATA ENGINEER WITH 5+ YEARS OF EXPERIENCE

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

PROFESSIONAL SUMMARY

Data Engineer with 6+ years building Snowflake and Spark pipelines for e-commerce and delivery businesses. Owns ingestion, modeling, and orchestration end to end, and works with analysts and product teams to keep core tables accurate and on time.

CORE SKILLS

Python | SQL | Snowflake | Apache Spark | Kafka | Airflow | dbt | AWS | Redshift | Great Expectations | Data modeling | Orchestration | CI/CD | Cost tuning | Data quality

PROFESSIONAL EXPERIENCE

Data Engineer Jun 2023 - Present
DoorDash, San Francisco, CA

- Own 45 Airflow pipelines feeding merchant and delivery tables in Snowflake, holding on-time landing above 99% over the past year.
- Moved order events from nightly batch to Kafka with Spark Structured Streaming, bringing dashboard freshness from a day to ten minutes.
- Modeled a merchant-performance mart in dbt with 120 tests, which analysts adopted in place of three older reporting tables.
- Trimmed Snowflake compute by 17% through clustering keys, warehouse sizing, and removing duplicate transforms.
- Created a data-contract check on the order event schema, blocking breaking field changes from reaching downstream tables.
- Mentored one new engineer through a first pipeline, reviewing the design and tests before it shipped to production.

Data Engineer Mar 2021 - May 2023
Chewy, Dania Beach, FL

- Built the ingestion framework for subscription order data from MySQL into S3 and Redshift using Python and AWS DMS.
- Added data-quality checks with Great Expectations to 30 jobs, catching late and duplicate files before morning reports ran.
- Led a migration of reporting jobs from Redshift to Snowflake with two engineers, finishing a week ahead of plan.
- Wrote on-call runbooks and joined the rotation, helping lower repeat pages on the pipeline team over two quarters.
- Cut a 4-hour customer-table rebuild to 35 minutes by switching it to incremental loads keyed on update timestamps.
- Partnered with analysts to retire 12 stale tables, after confirming through query logs that nothing downstream still read them.

Associate Data Engineer Jul 2020 - Feb 2021
Dell Technologies, Round Rock, TX

- Maintained ETL jobs in Informatica and SQL Server for supply chain reports, fixing recurring failures tied to late vendor files.
- Rewrote a slow inventory aggregation in Spark SQL, reducing the nightly run from three hours to about 50 minutes.
- Documented source-to-target mappings for 25 tables and trained two new teammates on the job schedule.
- Automated the monthly vendor-file validation with Python, removing a manual step that had taken one engineer half a day.
- Supported a data-center cost audit by exposing supply chain tables to finance through a governed, documented view.
- Joined weekly vendor calls to explain data requirements, which got two suppliers to send files in a fixed format and ended most late loads.

EDUCATION

B.S. in Computer Science, Texas A&M University, College Station, TX

CERTIFICATIONS

SnowPro Core Certification | Google Cloud Professional Data Engineer