

PW

Patricia Williams

Data Engineer

Senior data engineer with eight years across fintech and SaaS, focused on streaming systems and warehouse cost control. Lead by writing the first version of hard things and then handing them off.

Contact Information



(512) 555-0167



elena.vasquez@example.com



Austin, TX 12345

Education

M.S. Computer Science, University of Texas at Austin, 2016

B.S. Computer Science, University of California San Diego, 2014

Key Skills

- Python, SQL, Scala
- Apache Kafka and Flink for streaming pipelines
- Apache Iceberg, Parquet, and S3-based lakehouse design
- Snowflake and BigQuery warehouse architecture and cost optimization
- dbt modeling and data contracts
- Airflow DAG design and SLA monitoring
- Spark performance tuning
- AWS (S3, EKS, Lambda, IAM) and Terraform
- Event schema design and registry management
- CI/CD for data pipelines (GitHub Actions, dbt Cloud)
- Technical mentorship and hiring loop design

Professional Experience

SENIOR DATA ENGINEER I MAPLE RIDGE PAYMENTS, AUSTIN, TX 2021 – PRESENT

- Lead engineer on the merchant settlement pipeline: Kafka, Flink, Iceberg on S3, ~14M events per day.
- Reduced Snowflake spend from \$87K to \$52K per month by rewriting three credit-hungry models and switching warehouses to multi-cluster auto-suspend.
- Set up a data contracts process between platform and analytics; cut Slack escalations about broken schemas to roughly one per month.
- Run the data engineering hiring loop, including a take-home rubric I introduced after the old one was producing too many false positives.
- Mentor four engineers across two pods; two were promoted in the last review cycle.

DATA ENGINEER II I GREYLOCK TIDE SOFTWARE, SAN DIEGO, CA 2018 – 2021

- Built the company's first dbt project (180 models) and ran an internal SQL bootcamp for analysts.
- Migrated batch ETL from on-prem SQL Server to BigQuery over nine months without a customer-visible outage.
- Owned PagerDuty rotation for the data platform; reduced overnight pages by writing better DAG retry logic and SLA monitors.
- Designed the event schema registry used across three product teams.

DATA ENGINEER I GREYLOCK TIDE SOFTWARE, SAN DIEGO, CA 2016 – 2018

- Built the customer 360 mart that became the source of truth for marketing and finance.
- Wrote ingestion connectors for Salesforce, Zendesk, and Stripe in Python; later open-sourced two of them internally.
- Tuned a Spark job that was the team's slowest, taking it from 3.5 hours to 47 minutes.

- Incident response and on-call rotation leadership