



Edward Green

Data Engineer

Data engineer with six years building batch and streaming pipelines on AWS and GCP. Comfortable owning the path from Kafka topic to dbt model to BI dashboard, and known for cutting cloud bills by retiring pipelines nobody uses.

CONTACT INFORMATION



(503) 555-0142



priya.raman@example.com



Portland, OR 12345

EDUCATION

B.S. Computer Science, University of Washington, 2018

AWS Certified Data Analytics - Specialty, 2021

dbt Analytics Engineering Certification, 2022

KEY SKILLS

- Python (PySpark, Pandas, Airflow operators)
- SQL (Snowflake, BigQuery, Postgres)
- dbt (model design, tests, exposures, macros)
- Apache Kafka and Spark Structured Streaming
- AWS (S3, MWAA, Glue, EMR, Lambda) and GCP (BigQuery, Dataflow)
- Apache Iceberg and partitioned lakehouse design
- Airflow DAG authoring and migration
- Data quality tooling (Great Expectations, dbt tests)
- Terraform and CI/CD for data pipelines (GitHub Actions)
- Feature stores (Feast) and ML pipeline support
- Cost optimization and warehouse usage auditing

PROFESSIONAL EXPERIENCE

Senior Data Engineer | Northwind Forge Analytics, Portland, OR
2022 - Present

- Rebuilt the marketing attribution pipeline on Snowflake and dbt, cutting nightly runtime from 4 hours to 38 minutes.
- Migrated 60+ legacy Airflow DAGs from EC2 to MWAA; documented a runbook the rest of the team now uses for new jobs.
- Caught a join condition error in the subscription revenue model that was double-counting trial users, saving roughly \$410K in misreported ARR before the board deck went out.
- Mentor two junior engineers, including weekly code review sessions and a rotating on-call shadow program.
- Killed 14 unused Looker explores and 9 stale dbt models after a quarterly usage audit.

Data Engineer | Cobalt Harbor Logistics, Tacoma, WA
2019 - 2022

- Built a Kafka + Spark Structured Streaming pipeline ingesting GPS pings from 2,300 trucks at roughly 8K events per second.
- Designed partitioned Iceberg tables on S3 that replaced a Postgres warehouse straining at 4 TB.
- Wrote Great Expectations checks for the freight invoicing dataset; reduced finance reconciliation tickets by about a third.
- Partnered with the ML team to stand up a feature store on Feast for next-day delivery predictions.

Analytics Engineer | Cobalt Harbor Logistics, Tacoma, WA
2018 - 2019

- Owned the BigQuery layer for ops dashboards used by 40+ dispatchers across three regions.
- Converted ad-hoc SQL into a versioned dbt project, making model lineage visible for the first time.
- Drafted the team's first SQL style guide; still in use after two reorgs.

- Mentorship, code review, and on-call rotation design