Terrence Holcomb

New York, NY · terrence.holcomb@example.com · (212) 555-0148 · linkedin.com/in/example-terrence-holcomb

My Key Qualifications

- **11 years in data engineering** across fintech and payments: Betterment, Marqeta, and Green Dot, in one unbroken run since 2015.
- Owns batch and streaming architecture on **Kafka, Spark, Snowflake, and Airflow** over AWS, redesigning pipelines that cut end-to-end latency from 95 minutes to 18.
- Cut P1/P2 data incidents **from 11 per quarter to 4** by standing up SLA tracking, schema checks, lineage, and Datadog alerting across Airflow workflows.
- Payments and lending data domain hands-on: built **CDC-based ingestion** for authorization and settlement data at Marqeta, moving reconciliation reporting from next-day to under 30 minutes.
- Lowered Snowflake and S3 platform cost by **$34,000 per month** through partitioning, file sizing, and warehouse utilization work.
- Mentors data engineers through design reviews and incident response: supported a team of 5 and cut code review turnaround from 3 days to 1.
- Terraform, GitHub Actions, and CI/CD patterns in production: standardized orchestration raised on-time deployments **from 82% to 97%**.

Experience

Senior Data Engineer — **Betterment**, a digital investment and personal finance platform

Mar 2022 – Present · 4 yrs 8 mos

- Lead design of batch and streaming data architecture across AWS, Snowflake, Kafka, and Spark workflows, and set standards for data quality, schema management, lineage, and secure data handling in production pipelines
- Redesigned core transaction and account pipelines across batch and streaming workloads, combining Kafka event ingestion, Spark processing, and Snowflake serving models on AWS: reduced end-to-end data latency from 95 minutes to 18 minutes for 12 critical downstream datasets
- Cut P1 and P2 data incidents from 11 per quarter to 4 by implementing SLA tracking, schema checks, lineage documentation, and Datadog alerting across Airflow-managed workflows
- Lowered monthly Snowflake and S3 platform cost by $34,000 by reworking partition strategies, file sizing, incremental processing, and warehouse utilization patterns for finance and analytics workloads
- Led migration of legacy Python jobs into standardized orchestration and CI/CD patterns, raising on-time production deployments from 82% to 97% with reusable Airflow templates, GitHub Actions checks, and Terraform-managed infrastructure modules
- Mentored 5 data engineers through platform modernization, cutting median code review turnaround from 3 days to 1 day via design reviews, coding standards, and incident response runbooks

Sample resume made with Specific Resume · specificresume.com

Data Engineer — **Marqeta**, a card issuing and payments platform for embedded finance

Feb 2018 – Feb 2022 · 4 yrs

- Built CDC-based ingestion for payment authorization and settlement data into Snowflake, designing Kafka-backed ingestion patterns and SQL models for incremental processing; improved data availability for reconciliation reporting from next-day to under 30 minutes.
- Strengthened pipeline reliability for finance-close datasets by adding Airflow dependency controls, retry logic, and production alerting with Datadog, raising SLA attainment from 88% to 99.2%.
- Cut average Spark job runtimes on large transaction history backfills from 4.5 hours to 2.1 hours by optimizing joins, partitioning strategy, and storage layout on AWS-based processing jobs.
- Developed Python and SQL pipelines for payments, ledger, and reconciliation data, and maintained Airflow orchestration, warehouse models, and operational monitoring for downstream analytics.

Business Intelligence Engineer — **Green Dot Corporation**, a banking and payments technology company

Jan 2015 – Jan 2018 · 3 yrs

- Delivered dimensional models for card transaction and customer support reporting, designing fact and dimension tables and standardizing business logic in SQL, which enabled 9 recurring executive and operational dashboards.
- Improved daily ETL stability for reporting workflows by refactoring legacy SQL procedures and adding validation checks to batch processes, cutting failed jobs from 14 per month to 5.

Data Analyst — **Regions Bank**, a regional consumer and commercial bank

Nov 2013 – Dec 2014 · 1 yr 1 mo

Automated a recurring branch performance reporting process by replacing manual spreadsheet consolidation with scheduled SQL extracts and standardized report logic, saving 8 analyst hours per week.

Education

- **University of Alabama at Birmingham** — Bachelor of Science in Information Systems · 2009 – 2013

Skills

- **Data Engineering & Processing:** Python, SQL, Spark, Kafka, Apache Airflow
- **Cloud & Data Platforms:** AWS, Snowflake, S3, Datadog, Terraform
- **Data Architecture Methods:** CDC, Dimensional Modeling, Data Governance

Sample resume made with Specific Resume · specificresume.com