

PERSONAL INFORMATION

✉ marcus.wright@example.com

☎ +1 555 0121

📍 New York, NY

SKILLS

Core stack

- Kafka
- Flink
- Snowflake
- Kubernetes
- Terraform

Delivery

- Data modeling
- Data quality
- CI/CD
- Observability
- Documentation

Collaboration

- Requirements workshops
- Incident response
- Mentoring
- Agile delivery

EDUCATION

Carnegie Mellon University

Bachelor of Engineering

Computer Science and Data Systems

📅 09/2015 - 06/2019

LANGUAGES

English

Native

Mandarin

Professional working proficiency

CERTIFICATIONS

TOGAF Enterprise Architecture

Professional Training Board

📅 05/2024

Data Reliability Workshop

Data Engineering Guild

📅 08/2023

Marcus Wright

Principal Data Engineer

SUMMARY

Principal Data Engineer focused on global lakehouse architecture and governance. Builds reliable data products with Kafka, Flink, Snowflake, Kubernetes, Terraform, connects engineering decisions to measurable operations, and communicates clearly with cross-functional teams.

WORK EXPERIENCE

Principal Data Engineer

Cobalt Payments

New York, NY

📅 01/2022 - Present

- Designed Kafka and Flink pipelines for global lakehouse architecture and governance at Cobalt Payments; while partnering with product, platform, and finance leads, processing 3.8 billion payment events, and raised on-time delivery from 92% to 99.6%.
- Reworked Snowflake models and partitioning for the Global Settlement Lakehouse warehouse; coordinated a 9-person engineering group, cutting median query time from 24 seconds to 4.1 seconds.
- Instrumented schema, freshness, and lineage checks across 38 production tables for global lakehouse architecture and governance; worked with 14 analysts, reducing data incidents by 63%.
- Automated backfills, alerts, and runbooks around Kubernetes jobs in New York, NY; supported 12 markets and saved 86 on-call hours per quarter.

Analytics Engineer

Cobalt Payments Data Studio

New York, NY

📅 04/2019 - 12/2021

- Designed Kafka and Flink pipelines for global lakehouse architecture and governance at Cobalt Payments; with analysts, SREs, and domain owners across the delivery squad, processing 3.8 billion payment events, and raised on-time delivery from 92% to 99.6%.
- Reworked Snowflake models and partitioning for the Global Settlement Lakehouse warehouse; coordinated a 6-person engineering group, cutting median query time from 24 seconds to 4.1 seconds.
- Instrumented schema, freshness, and lineage checks across 38 production tables for global lakehouse architecture and governance; worked with 5 application teams, reducing data incidents by 63%.
- Automated backfills, alerts, and runbooks around Kubernetes jobs in New York, NY; supported 4 business units and saved 86 on-call hours per quarter.

PROJECTS

Global Settlement Lakehouse

Built a governed data product for global lakehouse architecture and governance, with documented ownership, tested transformations, and an analyst-ready serving layer.

Kafka Flink Snowflake Kubernetes