

Owen Fischer

Big Data Engineer

owen.fischer@example.com

+1 555 0115

Atlanta, GA

"Big Data Engineer focused on distributed network telemetry. Builds reliable data products with Kafka, Spark, Hadoop, Hive, Kubernetes, connects engineering decisions to measurable operations, and communicates clearly with cross-functional teams."

Education

Georgia Institute of Technology

Bachelor of Engineering · Computer Science and Data Systems

2015.09 - 2019.06

Skills

CORE STACK

Kafka Spark Hadoop

Hive Kubernetes

DELIVERY

Data modeling Data quality

CI/CD Observability

Documentation

COLLABORATION

Requirements workshops

Incident response Mentoring

Agile delivery

Work Experience

Big Data Engineer

SignalForge Telecom

2022.01 - Present

- Designed Kafka and Spark pipelines for distributed network telemetry at SignalForge Telecom; while partnering with product, platform, and finance leads, processing 8.4 PB of history, and raised on-time delivery from 92% to 99.6%.
- Reworked Hadoop models and partitioning for the Network Insight Grid 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 distributed network telemetry; worked with 14 analysts, reducing data incidents by 63%.
- Automated backfills, alerts, and runbooks around Hive jobs in Atlanta, GA; supported 12 markets and saved 86 on-call hours per quarter.

Languages

English Native

German Professional working proficiency

 Certifications

**Cloudera Data Platform
Generalist**

Professional Training Board
2024.05

Data Reliability Workshop

Data Engineering Guild
2023.08

Analytics Engineer

 SignalForge Telecom Data Studio

 2019.04 - 2021.12

- Designed Kafka and Spark pipelines for distributed network telemetry at SignalForge Telecom; with analysts, SREs, and domain owners across the delivery squad, processing 8.4 PB of history, and raised on-time delivery from 92% to 99.6%.
- Reworked Hadoop models and partitioning for the Network Insight Grid 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 distributed network telemetry; worked with 5 application teams, reducing data incidents by 63%.
- Automated backfills, alerts, and runbooks around Hive jobs in Atlanta, GA; supported 4 business units and saved 86 on-call hours per quarter.

Projects

Network Insight Grid	Built a governed data product for	
	distributed network telemetry, with	Kafka
	documented ownership, tested	Spark
	transformations, and an analyst-ready	Hadoop
	serving layer.	Hive