

# Aaron Ortega

aaortega18@gmail.com | github.com/aaron-ortega | linkedin.com/in/aortega04

---

## EXPERIENCE

---

### Senior Software Engineer — AgentSync

Jan 2023 – Present · Remote, TX

- Established CI/CD, automated testing, version control, and structured-logging standards for the data systems and mentored 5 engineers in adopting them
- Built two third-party data integrations (SureLC and DTCC) that flow insurance appointment contracts over SFTP and S3 into a centralized Go service that gates whether agents can legally sell, surfaced through the admin application
- Built the WebCE continuing-education integration on ECS and RDS Postgres, provisioned with Terraform, exposing two Go REST endpoints that validate state-by-state license-renewal compliance for thousands of agents, and shipped it as a standalone add-on customers purchase
- Wired three backend services onto a Kafka/Confluent event bus, integrating Salesforce through its Pub/Sub API to stream domain events past the REST API limits that had throttled high-volume usage
- Built internal Claude skills standardizing pull-request, code-review, and REST API conventions across 20 backend engineers, cutting recurring review feedback on error handling and API consistency

### Senior Data Engineer — Farmer's Fridge

Apr 2022 – Dec 2022 · Chicago, IL

- Introduced the Analytics team's first pull-request and CI/CD process (GitHub Actions), replacing ad-hoc deploys where engineers ran schema changes directly in prod from Slack messages, and reviewed roughly half of all PRs for query quality and consistency
- Refactored and parallelized a 400+ time-series workload onto Dask for distributed compute and Dagster for orchestration, cutting runtime by 62%

### Data Engineer II — Farmer's Fridge

Apr 2021 – Mar 2022 · Chicago, IL

- Owned observability for 600+ SQL models, wiring 1,000+ data-quality tests into CI to run on every pull request in non-prod, catching data issues before they reached production
- Migrated 20+ pipelines from cron to Dagster, adding automatic retries, structured logging, alerting, and simpler backfills to jobs that previously failed silently
- Integrated third-party sources (Salesforce, Stripe, and others) into Snowflake on scheduled Fivetran CDC syncs landing through S3, giving Analytics managed access to external data without custom ingestion code
- Rebuilt customer identity resolution in SQL, deduplicating customers for an accurate B2C unique-customer count and correct lifetime value, and removed its Python and AWS dependencies so the Analytics team could own it directly
- Implemented consumer-privacy data masking and tracking with dbt macros and Snowflake object tagging to meet California CCPA and CPRA requirements

### Data Engineer I (Founding Data Hire) — Farmer's Fridge

May 2020 – Mar 2021 · Chicago, IL

- Migrated a 1TB Aurora MySQL database to Snowflake in two months using Python and Apache Spark, centralizing it alongside the rest of the company's data for analytics and ML
- Streamed transactional MySQL changes in near real time through Kinesis to S3 and into Snowflake via Snowpipe, keeping the warehouse continuously in sync with production

## SKILLS

---

**Languages:** Python, Go, SQL **Tools:** Claude, Docker, dbt, Airflow, Dagster, Fivetran, Confluent, Datadog, Terraform, OpenAPI **AWS:** DynamoDB, Lambda, Redshift, S3, ECS, Kinesis **Data:** Snowflake, PostgreSQL, Spark, Kafka

## EDUCATION

---

**Master of Science in Physics** — University of Chicago