

# JORDAN ALEXANDER REYES

jordan@jordanreyes.dev | +14155550182 | linkedin.com/jordanreyes-sample

---

## SUMMARY

Senior Data & AI Engineer with 3+ years of depth in enterprise data engineering — lakehouse design, batch and streaming pipelines, governance — now building autonomous agent workflows and LLM-driven internal tooling on top of that foundation. Shipped subscriber-facing data products at scale for 18M+ users at Lumen Streaming, founded the data platform at Verdant BioSignals from zero, and at Northwind Logistics designed and piloted an autonomous incident-response agent projecting a reduction in pipeline downtime from 180 hrs/month to under 6.

## SKILLS

<b>Languages</b>	Java, TypeScript, Bash, JavaScript, Go (basic)
<b>Data Engineering</b>	Snowflake, Iceberg, Trino, PySpark, Apache Beam, Looker, Dagster, Segment, Apache Kafka, Data Modeling (Activity Schema, Inmon, SCD Type 4), Lakehouse Patterns, CDC Pipelines, Observability (Monte Carlo, Soda)
<b>AI Engineering</b>	Haystack, AutoGen, Weights & Biases, Vector DBs (Pinecone, Weaviate), Multi-Agent Orchestration, Hybrid Retrieval, Evals & Guardrails, Function Calling, Embeddings Tuning, Mistral API, Cohere API, Continue.dev
<b>Cloud &amp; DevOps</b>	Azure (Blob Storage, Synapse, Functions, Event Hubs, Key Vault, AKS, DevOps Pipelines, Container Registry, Service Bus), Vercel

## EXPERIENCE

### ***Northwind Logistics | Data and AI Engineer | May 2025 – Present***

- Designed and piloted an autonomous incident-response agent powered by LLMs — walking downstream lineage across Snowflake, Dagster, and Segment to independently identify, categorize, and explain pipeline failures for downstream teams and platform engineers; pilot projected to cut pipeline downtime from 180 hrs/month to under 6 hrs.
- Implemented the agent natively on Snowflake using Cortex Functions as invocable agent tools and Snowpark Container Services for hosting, with bespoke MCP servers bridging Dagster and Segment for lineage signals; published the deployed agent as a REST endpoint callable from Continue.dev, Cline, or any CI runner, with W&B Weave for full-trace observability.

### ***Verdant BioSignals | Lead Data Architect | Apr 2024 – Apr 2025***

- Stood up the entire data platform from greenfield — capturing, transforming, modeling, and serving 90+ metabolic markers, 280+ epigenetic signals, oral microbiome panels, and other longitudinal vitals to underpin a precision wellness program.
- Built a real-time event pipeline on Azure Event Hubs — ingesting 600K+ daily samples from always-on health feeds including ECG patches, sleep rings, smart scales, hydration trackers, and skin-conductance sensors — enabling near-instant clinical alerts and a deliberate transition from nightly batch to event-driven serving.

### ***Lumen Streaming | Data Engineer | Dec 2022 – Mar 2024***

- Spearheaded experimentation infrastructure for SkipAhead/Premium tiers, unlocking evidence-based product and bundling decisions across Growth and Subscriber teams.
- Built a transcript review system for Member Care (Lumen Voice Initiative) — automatically scoring whether agents adhered to Lumen's service playbook during member calls; rolled out company-wide for QA and coaching workflows.

## PUBLICATIONS

- Self-Documenting Data Products for Schema Evolution | Lumen DataConf '23

**EDUCATION:** Mid-State Polytechnic University | B.S. Computer Science | 2014 – 2018