

Matthew Nelson

Head of Data & Analytics · AI-native platform builder

Calgary, Canada · matthewpeternelson@gmail.com · +1-403-667-3022 · mnelson.ca · linkedin.com/in/matthewpeternelson

Data & analytics leader. I architect modern data platforms (Dagster, Snowflake, Polars, DuckDB) and ship AI-native products with Claude Code, MCP, and the AI SDK — 1.4B+ rows in production at StackDX, 16k+ products maintained by LLM-in-the-loop pipelines at Twochannel. P.Eng, Berkeley MIDS, 9+ years.

SELECTED IMPACT

Architected a production Dagster platform at StackDX across 15+ US states and 400+ transformation assets — 1.4B+ rows at the mart layer, 400M+ in downstream exposures.

Cut critical-dataset processing time via targeted Polars and DuckDB rewrites, incremental materializations, and coalesced fact/mart layers.

Founded twochannel.ai: an AI-native monorepo where LLM-in-the-loop pipelines and a custom Claude Code toolchain maintain 16k+ products across 2k+ brands with minimal manual work.

Owned Canlin's annual NI 51-101 corporate reserves for 5 years; technical revisions added \$730M NPV10 over 4 years.

Built the dbt Cloud / Snowflake platform that became Canlin's foundation for operational data, BI, and ML; automation saved 50+ hours/week across Operations.

SKILLS

AI Engineering & Agentic Development: Claude Code (agents, skills, hooks, slash commands, sub-agents), MCP servers & clients, Anthropic / OpenAI / Perplexity APIs, Vercel AI SDK, Structured outputs & schema-adherent LLM calls, Prompt caching & evals, RAG, Agent frameworks (autogen, crew), LLM-in-the-loop data pipelines, Editors: Claude Code, Cursor, Windsurf

Data Engineering & Modern Data Stack: dbt (Core & Cloud), Dagster, Snowflake, Snowpark, Polars, DuckDB, Tableau Prep, cube.js & dbt semantic layers, pandas, SQL, Schema contracts

Cloud, Infra & DevOps: AWS (ECS, RDS, S3), Terraform, Docker, GitHub Actions CI/CD, Vercel, Railway, Supabase, Doppler, pre-commit / ruff / pyright

Machine Learning & NLP: scikit-learn, XGBoost, PyTorch, TensorFlow / Keras, Hugging Face, nltk, Time-series forecasting, Classification, Anomaly detection

Web App Development: Next.js, React, Astro, Directus (headless CMS), Algolia, Node.js, Express, Tailwind CSS, Playwright, Streamlit, Flask, Dash, Turborepo / npm workspaces

Programming Languages: Python, SQL, TypeScript, JavaScript, Bash, R

EXPERIENCE

Data Platform Engineer · Stack Technologies Ltd.

Jun 2025 — Present

Lead engineer on the USA data platform rebuild — replacing a legacy Python ETL stack with a production-grade Dagster-native architecture on AWS, with Polars and DuckDB doing the heavy transformation work. Operating at scale: 1.4B+ rows at the mart layer, 400M+ rows in downstream exposures, ~47GB of parquet across landing/intermediate/marts. Operating as an AI-native developer: custom Claude Code skills, sub-agents and hooks act as force multipliers across scaffolding, data inspection, formatting and review — measurably increasing the rate at which new state pipelines and transformation assets ship.

Architected production Dagster platform across 15+ US states and 400+ transformation assets (1.4B+ mart rows, 400M+ exposure rows), replacing a legacy Python ETL stack with a schema-contracted, testable modern data platform.

Cut processing time on critical datasets through targeted Polars/DuckDB rewrites, incremental materializations, and coalesced fact/mart layers.

Built a project-specific Claude Code toolchain — sub-agent, slash commands, auto-format hook — that measurably increases the rate of shipping new pipelines.

Founder & Principal Engineer · Twochannel

Jan 2021 — Present

Designed and built Twochannel end-to-end: product catalog, ingestion pipelines, search, recommendations, and frontend. The catalog — 16k+ products across 2k+ brands — is maintained almost entirely through AI-native workflows: LLM-in-the-loop scrapers, structured-output enrichment, semantic dedup, and a Claude Code toolchain operating against a full monorepo. A working proof that a single engineer with modern AI tooling can run the data engineering of a team.

Brand-scrapers → LLM normalizer → reviewed dedup pipeline now maintains 16k+ products across 2k+ brands in Directus with minimal manual intervention.

Claude Code is the primary IDE for Twochannel: skills, sub-agents, hooks, and MCP servers do the undifferentiated heavy lifting; human focus stays on product decisions.

Turborepo monorepo shipping Next.js (Vercel), Directus (Railway), Algolia search, Neo4j relationships, and Doppler-managed secrets across three hosting targets.

Senior Data Analytics Engineer · Paramount Resources Limited

Apr 2024 — May 2025

Senior member of the Data Analytics and Integration team — integrating best practices across the full data stack, mentoring ICs, and advocating for a modern data platform. Delivered polished, robust solutions spanning custom web apps, modern data-stack PoCs, data engineering pipelines, and high-fidelity dashboards. Toolkit: Databricks, PowerBI, Spotfire, dbt Core, cube.js, DuckDB; SQL, Python, R, DAX.

Multi-level Netback & Operating Cost PowerBI dashboard with embedded planning/forecasting, from corporate view to well level.

Streamlit app replacing manual zipper-frac sequencing with an algorithm optimizing surface efficiency, frac azimuth, and crane-bank rotation.

Local modern data stack PoC (Python / dbt / cube.js / DuckDB / PowerBI in Docker) replacing legacy R + CSV processes.

Data & Advanced Analytics Lead / Corporate Data Scientist & Reserves Manager · Canlin

Oct 2017 — Apr 2024

Energy Corporation

Designed, built, and ran Canlin's industry-leading data stack (dbt Cloud, Snowflake, DataRobot, Tableau, Streamlit). Owned the annual NI 51-101 reserves process and drove +\$730M NPV10 through technical revisions. Mentored juniors into data roles.

Built the dbt Cloud / Snowflake platform that became the foundation for all operational data, BI, and advanced analytics at Canlin.

XGBoost 30-day production forecaster for 8,000+ wells, inferring daily into corporate reporting.

Automated weekly production report saving 50+ hours/week across Operations.

Exploitation / Development Engineer · Centrica Energy Canada (Jul 2014 — Sep 2017) — Gas-plant longevity analysis setting 5–10-year corporate direction for Foothills assets.

Production & Exploitation Engineer · Pengrowth Energy (2010 — 2014) — On-site SAGD production engineering at Lindbergh plus drill-program design in Southern Alberta.

EDUCATION

Master of Information and Data Science · University of California, Berkeley (May 2016 — Jun 2018, GPA 3.888)

Bachelor of Applied Science, Chemical Engineering · University of Waterloo (Sep 2005 — Jun 2010)

Professional Engineering Designation (P.Eng) · Association of Professional Engineers and Geoscientists of Alberta (APEGA) (2014)