Job description: Data Intelligence Engineer

The role. A data team has one job: enable the business to make good decisions, faster, based on a more accurate picture of what's happening. Agents now write most of the code. You create and conduct them, and you're accountable for what they build: when a number matters, you're the one who signs off and says "yes, this is right." Do it well, and everyone who asks a question gets the experience of working with the smartest analyst in the company.

The goal is self-service. Anyone in the organization, including people without the word "Data" in their title, should be able to open Claude, ChatGPT or another agent, answer questions about the business correctly, build apps and visualizations, and make high-quality decisions, asking the follow-up while the context is still in their heads.

What you'll do:

- **Conduct the agents.** Automate ingestion and transformation with agents, decide what's ready for production, and own the platform.
- **Build a clean data model.** It comes first. A simple, clear model saves agents a lot of time and energy (and tokens).
- **Curate the context.** Write down what the model can't make obvious: which fields join, what "ARR" means, why sales and finance count revenue differently. Keep a clear ontology. Context is a manual for agents; the cleaner the model, the shorter the manual.
- **Run evals.** They're system tests: questions with known answers, run again and again, so you know the answers are right and find out when something breaks.
- **Supervise in proportion to the blast radius.** A Wild West, where anyone builds on a reasonable foundation and you keep your hands off. A Curated set of the dashboards, apps and metrics that matter, which you make a best effort to fix. And the Core: key metrics, high quality and predictable at all times. The board number lives in the Core.
- **Keep it working.** Build frameworks, prevent breakage, fix what breaks anyway, and be accountable for the surfaces the business depends on.

You might be a fit. This is a role you play, whatever your title says. It fits if you:

- are a data engineer, an analytics engineer, or an analyst, and have watched those roles blend
- already use agents for SQL and pipelines, and know where they go wrong
- would rather fix the data model than write a manual for a messy one
- believe data engineering is an accountability function, regardless of tooling or who does it

How you'll know it's working. More people answer their own questions, come back with the next one, and make confident decisions faster. When an answer drifts, your evals catch it first. You have no backlog and more to do than ever.

