

Enterprise-Grade Analytics Without the Enterprise Bill

THE CHALLENGE

A 200-employee textile manufacturer had outgrown spreadsheets and was trying to get live reporting directly off their operational PostgreSQL database. It wasn't working, running analytical queries against a live transactional system slowed down the actual business software people depended on every day. Pricing out enterprise data platforms revealed thousands of dollars a month in cost, and infrastructure that would strain their existing legacy systems rather than work with them.

OUR APPROACH

We kept managed PostgreSQL as their transactional database, exactly what it's built for, but separated the analytics layer out entirely so heavy reporting queries never touch the live system again. For transformation, we used DuckDB, which runs in-memory on serverless infrastructure with no standing server cost, paired with dbt to write clean, modular, version-controlled SQL.

Google BigQuery still has a role, but not as the everyday engine, strictly long-term cold storage, with a monthly archive push, so historical data stays searchable without paying to keep a massive warehouse running 24/7. Because the transformation logic lives in dbt, it's self-documenting, no single-developer "hostage" risk.

THE RESULT

A full modern data stack running for roughly \$60 to \$230 a month, against enterprise-equivalent pricing between \$1,200 and \$3,000+ a month for the same job.