Jay Bennett

AI Solutions Engineer | LLM Applications · Data Platforms

Minneapolis, MN | 651-983-0103 | jaybenne@gmail.com | linkedin.com/in/jay-b-301863a

SUMMARY

AI solutions engineer with 15+ years of enterprise data engineering behind the AI work. I build working LLM applications: a multi-tenant document platform with Claude-based classification and row-level security, a multi-agent analysis app, and a dealer-email parser that feeds a Snowflake warehouse. I focus on guardrails, validation, and human review so AI output is safe to use in regulated settings, and on getting the data underneath it right.

CORE SKILLS

AI / LLM: Claude & OpenAI APIs, multi-agent orchestration, semantic retrieval, LLM document extraction and classification, prompt and guardrail design, AI-assisted development (Claude Code, GitHub Copilot)

Application: React/TypeScript, Node.js, Python, Supabase (Postgres, RLS, Edge Functions), AWS Lambda/Amplify, Playwright, CI/CD

Data: Snowflake, Databricks, SQL Server, PostgreSQL, ETL, metadata-driven automation

Certifications: Advanced Prompt Engineering & Generative AI (Deloitte), Azure AI-900, SnowPro Core, AWS Cloud Practitioner

EXPERIENCE

Independent AI Development | Consultant & Builder — Jan 2026 – Present

- Dealer run parser: turns fixed-income dealer emails into a normalized daily residential mortgage-backed securities (RMBS) spread grid. Three-layer Snowflake warehouse (RAW, STG, MART), regex-first parsing with a large language model (LLM) handling the leftovers, and a FINRA API backfill. A 53-test suite caught eight bugs before release
- Legacy modernization tools: an Excel Logic Extractor (React, Claude API, SheetJS) that pulls embedded business logic out of workbooks, and a SQL stored procedure debugger demo with explain, optimize, and dependency actions. Both parse files directly with no server connection
- NBA officiating analytics engine: built a 53-game play-by-play corpus and tested 18 variables with a time-based train/test split. Reported the null results honestly (16 of 18 showed no signal) along with the two that held up
- Council: multi-agent web app where five specialist agents analyze a question in parallel and a judge agent reconciles them. Next.js and Supabase with authentication and usage tiers
- Divorce tax planning engine: browser-based federal and Minnesota 2026 tax engine with child support calculator, scenario comparison, Claude-based document extraction, and 19 test profiles validated against the original Excel/VBA workbook

Stack: Snowflake, Python, React, Next.js, Supabase, Claude API, AWS (Lambda, API Gateway, DynamoDB, S3, CloudFront), Vercel, Swift/SwiftUI

xBrezzo, LLC | Independent Consultant — Aug 2025 – Dec 2025

- Built a multi-tenant SaaS platform that lets tax practitioners collect, classify, and organize sensitive client financial documents
- Designed a Claude API document pipeline that classifies uploads and renames them to a standard Institution_AccountType_Name_Date convention
- Implemented row-level security (RLS), four-role access control, and multi-tenant data isolation on Supabase/PostgreSQL, with automated Playwright tests for each security boundary
- Built Supabase Edge Functions for case-summary generation; moved from OpenAI to Claude API after hands-on evaluation

Stack: Supabase (Postgres, RLS, Edge Functions), Claude API, OpenAI API, React, Playwright, GitHub, Claude Code

Deloitte | Senior Consultant — Jun 2023 – May 2025

- TN CCWIS Modernization (Snowflake): built Snowflake data models and ETL pipelines supporting Tennessee's child-welfare migration from a legacy platform to Salesforce
- Wrote Python automation that parses Salesforce metadata and report specs to generate data warehouse mappings, cutting manual modeling effort about 70% across four delivery tracks and producing 1,000+ Snowflake tables
- Created metadata-driven rules for fact/dimension classification, schema-change detection, and automated Snowflake schema updates

- FedEx DataWorks (Databricks): migrated Oracle workloads to Azure Databricks, implementing medallion-architecture transformations in Spark SQL and Python with validation frameworks and PL/SQL optimization
- Participated in Deloitte's AI Ambassador program (Microsoft Fabric and Copilot)

Stack: Snowflake, Databricks, Spark SQL, Python, Salesforce, Azure, Terraform, PL/SQL, SSIS, Power BI

Augeo | SQL Developer — Mar 2022 – Jul 2023

- Built and maintained SSIS ETL pipelines, stored procedures, and ad-hoc queries for participant and performance data across client programs
- Migrated on-premises SQL Server workloads to AWS RDS/EC2, owning data loads and validation

Stack: SQL Server, SSIS, AWS RDS/EC2, Azure DevOps

1upHealth | Data Engineer (Contract) — Apr 2021 – Jul 2021

- Onboarded new clients' FHIR healthcare data to meet a July 1 compliance deadline, configuring Apache NiFi flows to convert and load records

Stack: AWS, Terraform, PostgreSQL, Docker, Apache NiFi

USCG LLC | Cloud Engineer (Contract) — Oct 2020 – Mar 2021

- Designed cloud deployment pipelines for Python Flask web applications on AWS and GCP (Docker, NGINX, uWSGI) and introduced GitHub-based version control

EARLIER EXPERIENCE

Capella University, BI Developer (2014–2015) · Target, .NET Developer (2013–2014) · Rust Consulting, SQL Server Developer (2012–2013) · Best Buy, .NET Developer (2011–2012)

U.S. Bank, Systems Analyst / Programming Analyst / Bond Operations (2006–2011): built Access/VBA and VB.NET applications that automated escheatment and bondholder operations

EDUCATION & CERTIFICATIONS

Professional Certificate, Applied Data Science, Thayer School of Engineering at Dartmouth (2021)

B.A., Psychology & Business, University of St. Thomas

Certifications: Snowflake SnowPro Core · Power BI Data Analyst (PL-300) · AWS Cloud Practitioner · Azure AI-900, AZ-900, DP-900 · Advanced Prompt Engineering & Generative AI (Deloitte)