

Arthi Arumugam

Salem, India · +91 96776 67794 · arthi1805@gmail.com · linkedin.com/in/arthiarumugam99 · github.com/arthi-arumugam-git

Forward deployed engineer with 3+ years building production AI applications with Claude and deploying them inside enterprise customer environments. I have built with Anthropic's models specifically — Claude API powers Gnani, an open-source agentic spreadsheet with a full autonomous tool loop, and I use Claude as my primary AI engineering platform across all production work. I care deeply about AI being safe and beneficial — not as a talking point but because I have thought seriously about what it means to deploy AI systems that real people depend on. I operate autonomously under ambiguity, have delivered production AI platforms for Walgreens (\$5M impact), Grab, and Lumen Technologies, and am ready to represent Anthropic at the highest level in enterprise customer environments.

TECHNICAL SKILLS

Claude and LLMs	Claude API (Anthropic) — primary AI engineering platform. Advanced prompt engineering, structured output, LLM tool-use, multi-agent orchestration, agent skills, MCP server patterns, sub-agent design, evaluation frameworks, deployment at scale
Agentic systems	Compound AI workflow design, autonomous tool loops, agentic orchestration, output validation, graceful failure handling, Vercel AI SDK, LangChain patterns, RAG pipeline design, vector embeddings
Engineering	Python, TypeScript, Node.js, Next.js, REST APIs, Docker, GitHub Actions, n8n, Inngest, Supabase, SSE streaming, WebSocket, Scikit-learn, pandas
Cloud and data	GCP (BigQuery, Vertex AI), Azure (Synapse, Data Factory, OneLake), Microsoft Fabric, Databricks, Snowflake, Oracle PL/SQL, PostgreSQL, MySQL
Enterprise	SAP integration, Oracle ERP, IoT data pipelines, enterprise IT systems, regulatory compliance workflows, audit documentation, production deployment patterns
Certifications	Google Cloud Professional Data Engineer · Azure Data Engineer Associate · Microsoft Fabric Data Engineer Associate · Power BI Data Analyst Associate

EXPERIENCE

Informatics consultant, data and cloud engineering

Lumen Technologies · Remote · Oct 2025 – present

- Production GenAI agent: built a Vertex AI (Gemini)-powered BigQuery cost intelligence agent — a production system that analyses live query patterns through an agentic prompt loop and generates SQL optimisation recommendations. Designed the tool loop, handled edge cases where the model returned unexpected outputs, validated recommendations before surfacing them to engineers. This is real LLM deployment work — not a demo — and the same pattern I would apply building with Claude in a customer environment.
- 94% cost reduction: the agent's recommendations drove a BigQuery partitioning and clustering strategy cutting bytes scanned from 565 MB to 35 MB per query. Projected \$4,500 in annual savings per production table. Turned AI output into measurable production impact.
- Ambiguous, high-stakes environment: joined mid-transition with no handoff, no documentation, and active regulatory deadlines across 8 US states. Reconstructed the PUC filing process from scratch, filed on time, and built SOPs so the knowledge survives. This is the Anthropic FDE operating context: autonomous, no playbook, high stakes, must deliver.
- 90x faster reporting: automated the Oregon daily trouble report from 45 minutes to 30 seconds. Identified the pain, built the solution, adopted team-wide. Demonstrates bias for action.

Data analyst, AI and cloud platform engineering

Codesizzler · Remote · Jan 2024 – Oct 2025

- \$5M revenue protected — Walgreens: owned end-to-end delivery of a production AI supply chain platform across 200+ stores. Ran discovery workshops with Walgreens' engineering and ops teams, designed the

integration architecture (SAP + Oracle + IoT data into Microsoft Fabric OneLake), built anomaly detection models, deployed to production, and trained the customer team for independent operation. Stayed until the outcomes were real. This is the Anthropic FDE model: embed, understand, build, ship.

- +34% forecast accuracy: built predictive models with customer feedback loops. Evaluated outputs against historical ground truth before presenting to stakeholders. Did not ship until the numbers were real.
- Real-time pipeline — Grab: designed Azure Data Factory + Synapse pipeline from raw IoT event streams to live dashboards across Southeast Asian markets. Led architecture reviews with Grab's engineering team. Translated Grab's operational context into technical decisions — cross-functional at every level.
- \$42K savings — ABARTA: delivered enterprise BI platform across 7 business units. Built playbooks and documentation for customer self-sufficiency post-handoff.

Project management analyst

Network Rail · Three Bridges, UK · Apr 2023 – Jan 2024

- Power BI dashboards surfacing 12 safety issues invisible to leadership. SharePoint + Power Automate workflows reducing manual reporting by 30%. Documentation adopted across the operations team.

OPEN-SOURCE PROJECTS

Gnani · github.com/arthi-arumugam-git/gnani

Agentic spreadsheet — Claude API, autonomous tool loop, MCP-compatible patterns, Apache 2.0

- Built a full-stack agentic spreadsheet where =AGENT('prompt') triggers a Claude-powered agent that autonomously calls APIs, scrapes web pages, writes values across cells, applies formatting, and renders charts — without further human input after the prompt. Built specifically with Claude API because Claude's tool-use and structured output capabilities made the agent more reliable than alternatives I tested.
- Designed the compound AI architecture from scratch: tool schemas, output validation before writing, partial failure handling, and graceful degradation. The design mirrors Anthropic FDE deliverables: MCP server patterns, sub-agent tool loops, and agent skills for production workflows.
- Apache 2.0. Docker support. Self-hosting documentation written for other developers. Other engineers run it in production — the documentation had to be real.

Costctl · github.com/arthi-arumugam-git/costctl

Database cost optimisation CLI — Python, 6 enterprise platforms, 301 tests, GitHub Action

- Python CLI autonomously scanning BigQuery, Snowflake, Databricks, PostgreSQL, Redshift, and MySQL for cost waste. Each finding includes estimated monthly savings and a ready-to-run DDL fix. Ships as a GitHub Action. 301 tests. Runs in production pipelines without supervision.

OTHER PROJECTS

TaskMelt · taskmelt.app · [iOS App Store](#)

Voice-to-task AI app — Deepgram real-time STT, RevenueCat billing, shipped to App Store

- Built and shipped an iOS voice-to-task app using Deepgram real-time streaming STT API. Designed the full audio pipeline: microphone capture, live transcription, LLM task extraction, structured output. RevenueCat subscription billing. Solo-built from zero to App Store.

Lore · lore.surf

MCP server — connects Notion, Confluence, Slack, GitHub, Linear into AI coding tools

- Built an MCP server connecting organisational knowledge into Cursor and Windsurf. Demonstrates direct experience with the MCP protocol that Anthropic FDEs deliver as customer artifacts.

EDUCATION

M.Sc. Management (business specialisation)

Queen Mary University of London · 2022 · GPA 3.3 · School of Business Scholarship.

B.Tech Information Technology

St. Joseph's College of Engineering, Anna University · 2020 · 8.23/10 · Cyberbullying detection using NLP and Naive Bayes (87% accuracy).

