Karthick Govindaraju

Seattle, WA | karthickmit25@gmail.com | karthickg.dev | linkedin.com/in/karthickgovin | github.com/kgcodes1288

Forward Deployed / Applied AI Engineer | LLM Applications | Data Engineering

PROFESSIONAL SUMMARY

Data leader with 15 years of experience building production data systems, now working hands-on as an applied AI engineer. Builds end-to-end LLM applications at work and in personal projects (agents, MCP servers, LLM-in-the-loop pipelines, and agentic applications on Amazon Bedrock AgentCore and Anthropic Claude), backed by full-stack engineering (React, Node.js, Python, PostgreSQL, Docker) and large-scale data engineering (Redshift, PySpark/EMR, Airflow). At Amazon Ads, lead a BI engineering team and delivered an agentic AI application that saves 500+ hours per week across 50 users. Seeking a Forward Deployed Engineer, Applied AI Engineer, or GenAI Solutions Architect role.

CORE SKILLS

AI / LLM: Amazon Bedrock AgentCore, Bedrock Knowledge Bases, Anthropic Claude, Amazon Bedrock, LLM agents & tool use, Model Context Protocol (MCP) servers, prompt engineering, RAG/context grounding, evals & guardrails, Whisper speech-to-text, text-to-speech, LLM-based content moderation

Languages: Python, SQL, JavaScript/Node.js, HTML/CSS, Bash

Application Stack: React, Vite, Express, Flask, REST APIs, OAuth 2.0 / JWT / SAML SSO, PostgreSQL, Prisma, Supabase, SQLite, Twilio, webhooks

Infrastructure: Docker / Docker Compose, Linux, Caddy, Cloudflare Tunnel & Zero Trust, AWS (EC2, ECS Fargate, Lambda, S3, Glue, Athena), CI/CD, staging/prod environments, backups & observability

Data Engineering: Amazon Redshift, PySpark on EMR, Apache Airflow, ETL pipelines, data modeling, Pandas, Tableau, QuickSight

Customer-Facing: Requirements discovery, solution design, stakeholder demos, technical documentation, vendor integration, production support

PROFESSIONAL EXPERIENCE

Manager, Business Intelligence, Ads Sales (Twitch Ads, Agency, ADSP, Partner Dev) **08/2021 - Present**

Amazon Ads, Seattle, WA

- Owned product and solution design for an agentic AI application on Amazon Bedrock AgentCore: a web front end dispatches questions and tasks to agents running in isolated AgentCore Runtime sessions, which query datasets, run computations and reconciliations, and generate CSV, PDF, or HTML dashboard outputs. Designed the tool layer (custom agent skills plus Bedrock Knowledge Bases) and least-privilege IAM role-based access to datasets. In production with 50 finance and analytics users, saving each 10+ hours per week (500+ hours weekly) of reporting and reconciliation work.
- Delivered an earlier LLM analytics chatbot (Bedrock, Anthropic Claude 3, Python) grounded in curated revenue datasets, with answers validated against source-of-truth dashboards; replaced manual SQL requests from sales leadership.
- Re-architected legacy Redshift SQL ETL to PySpark on AWS EMR, cutting job runtime 50% and saving $50K per month ($600K/yr) in compute cost.
- Led the $2MM Tableau Server platform launch for the Ads Finance org: a 4-node cluster on AWS EC2 in a corporate VPC with SAML SSO and Active Directory permissions; scaled to 200+ dashboards serving 5,000+ users across Ads Sales and Finance.
- Led a critical data migration of Advertiser/Agency mappings across 50K advertisers and 1,000+ partners, covering data reconciliation, gap analysis, and revenue-impact assessment; drove the product fixes that closed every gap with zero reporting disruption.
- Lead a team of 6 BI engineers owning revenue reporting, data pipelines, and data marts for Twitch Ads, Agency, and Ad-Tech segments, serving global Ads Sales, Finance, and Compensation teams across multiple countries; own hiring, mentoring, roadmap, and prioritization.
- Owned vendor selection, contract negotiation, technical requirements, and integration for third-party ad-market insights data, enabling Sales to measure market share and set data-driven sales targets.

Data Engineering Lead **01/2016 - 08/2021**

Gulfstream Aerospace, Savannah, GA

- Built a web-based recommendation engine (Python Flask; market-basket analysis + Levenshtein text similarity) over 1M nonconformance records that surfaced historical fixes for current engineering issues, saving engineers 2+ hours/day and cutting solution lead time 80%+.
- Led end-to-end data infrastructure for the Engineering organization, including an Amazon Redshift data warehouse, dimensional data models (star/snowflake), Python/SQL ETL pipelines, and customer-facing Power BI dashboards.
- Delivered business-process automation projects used by 1,000+ engineers, saving $5M+ annually.

Technical Program Manager / Operations Engineer **06/2011 - 12/2015**

Gulfstream Aerospace, Savannah, GA

- Secured $1M in savings through Lean Six Sigma process-improvement initiatives, using hypothesis testing and design of experiments (Minitab) for root-cause elimination.

PERSONAL AI PROJECTS

CleanStay: Property-Operations Web App with AI Assistant Integration **2025 - Present**

Personal project | github.com/kgcodes1288/KASA

- Built an MCP (Model Context Protocol) server with OAuth 2.0 authorization so users can connect AI assistants such as Claude to their account and manage listings and jobs in natural language.
- Designed and built a full-stack property-operations web app for short-term-rental hosts, with Airbnb iCal sync that auto-generates cleaning jobs, room-by-room checklists, recurring maintenance, a vendor quote and owner approval flow that auto-drafts invoices, and role-based owner, co-host, and contractor dashboards.
- Built the full stack: React 18 + Vite, Node.js/Express REST API, PostgreSQL + Prisma, JWT and Google OAuth 2.0, Twilio SMS, Resend email, Cloudinary; Dockerized behind Caddy with isolated staging and production environments and one-command deploys.
- Ran phased end-to-end QA with synthetic accounts across UI, email, and SMS flows; diagnosed and fixed a deployment incident (container DNS aliasing routed prod traffic to staging) and hardened environment isolation.

Home Lab: Self-Hosted AI Automation Platform **2025 - Present**

Personal project

- Drive day-to-day automation with agentic AI (Claude Code + MCP servers for email, Google Workspace, and custom APIs); e.g., migrated ~50K photos/videos (404 GB), removed 12,158 duplicates, and auto-built albums from date + GPS clustering (+3,522 assets) with full undo logs.
- Architected and run a home-lab platform running 30+ containerized services (Docker Compose; Supabase/Postgres, Nextcloud, Immich ML, Jellyfin, multiple web apps) with Cloudflare Tunnel + Zero Trust Access, nightly encrypted restic backups to Backblaze B2, UPS-triggered graceful shutdown, systemd timers, and push alerting.

Curious Kutties: AI-Assisted Video Production Pipeline **2026 - Present**

Personal / family hobby project (kids' science videos)

- Built an end-to-end AI-assisted video production pipeline in Python that transcribes narration with Whisper, has an LLM generate scene plans (JSON), pulls matching stock footage from the Pixabay API with visual review sheets, renders 1080p episodes with titles, quiz cards, and transitions, mixes voice and music with automatic ducking, and publishes through the YouTube Data API.
- Built a community idea board with voting and a three-layer moderation pipeline: a rules engine, then LLM vetting for kid-safety and relevance (English, Tamil, Hindi), then human-in-the-loop admin review behind Cloudflare Zero Trust SSO.

EDUCATION & CERTIFICATIONS

M.S., Computational Data Analytics, Georgia Institute of Technology 2021
M.S., Industrial Engineering, University at Buffalo, SUNY 2011
B.E., Electronics & Instrumentation, Madras Institute of Technology, India 2009

Certified Six Sigma Black Belt (CSSBB), University at Buffalo | MicroMasters in Analytics, Georgia Tech