

RAMYA AUTHAPPAN

Agentic AI Engineer | Applied LLM Engineering

Chennai, Tamil Nadu, India | +91-9500073277 | ramyaat2025@gmail.com
linkedin.com/in/ramyaat | ramya-site.vercel.app | github.com/ramyaat

SUMMARY

Agentic AI Engineer with 18+ years shipping production systems for FAANG, FinTech, and high-growth SaaS (Amazon, PayPal, GitLab, Freshworks, Circle), and hands-on experience building agentic AI systems using the Claude Agent SDK, Claude Code MCP, LangChain/LangGraph, and RAG pipelines. Most recently deployed at Circle, building an autonomous, self-healing test agent that reads PR diffs, reasons about blast radius, and reduces CI cost by 50% — the same skill set (production integration, LLM orchestration, client-facing delivery) Agentic AI Engineer roles require. Combines deep systems/DevOps expertise with a verification mindset: "I build AI agents, and I know how to prove they work."

CORE TECHNICAL SKILLS

Agentic AI / LLM Engineering: Claude Agent SDK, Claude Code MCP, Buildkite MCP, Model Context Protocol (MCP), LangChain, LangGraph, multi-agent orchestration, RAG pipeline design, LLM-based root-cause analysis, prompt engineering, LLM routing (LiteLLM, multi-tier model hierarchies), self-healing / autonomous test agents, Amazon Bedrock AgentCore, LangSmith (agent tracing/observability), OpenRouter

Languages: Python, JavaScript, TypeScript, Ruby, Java, SQL

Cloud, DevOps & Integration: Docker, Kubernetes, GitLab CI, Jenkins, Buildkite, GitHub Actions, REST APIs, Node.js/Express, PostgreSQL, WhatsApp Cloud API, Webhooks, n8n workflow automation, self-hosted VPS deployment

Full-Stack / Product Engineering: Next.js 15, React 19, Tailwind CSS, shadcn/ui, FastAPI, Amazon CloudWatch

AI Evaluation & Reliability: PR-diff impact analysis, selective test execution, flaky-test triage, test observability, SLOs/SLIs, quality gates, evaluation-driven development, reference-dataset curation, model-eval vs. product-eval design

Client-Facing & Delivery: Stakeholder mapping, requirements discovery, RAPID decision documentation, technical enablement, cross-functional collaboration, engineering management, public speaking

APPLIED AI / AGENTIC AI PROJECTS (FDE ACADEMY)

ShipRight Logistics — Client Discovery & Solution Design Case Study

Jul 2026

- Produced stakeholder mapping, problem definition, and RAPID decision documentation for a fictional enterprise logistics client, simulating the discovery-to-recommendation workflow used in real FDE engagements.

Meeting Action-Item → GitHub Issue Agent (LangChain)

Jul 2026

- Built a LangChain agent that reads a meeting transcript, classifies each action item as confirmed, tentative, or a suggestion, and files a GitHub issue only for confirmed commitments — skipping tentative items and mere suggestions.
- Designed a single atomic tool call (create-issue + database update together) to eliminate a race condition where the model could batch tool calls out of order; the agent loops until resolved, then emails a summary table of created and skipped items via Gmail SMTP.
- Used Gemini 2.5 Flash via langchain-google-genai, with SQLite for state tracking and the GitHub REST API for issue creation.

Tech: Python, LangChain (LCEL), Gemini 2.5 Flash, SQLite, GitHub REST API, Gmail SMTP, Redis

SetuHaul — Agentic Freight Dock Scheduling Console (FDE Academy)

Aug 2026

- Built a driver-facing dispatch console for freight dock scheduling that replaces a rules engine with an LLM agent for operational reasoning — turning free-text driver reports of delays and dock-availability questions into grounded scheduling answers instead of plausible-sounding guesses.
- Constrained the LLM to a single tool that calls back into the backend over HTTP to check real, currently-open slots, so it can report availability but never invent it; booking itself is a separate deterministic REST call with no LLM in the loop, keeping reasoning AI-assisted while state changes stay fully deterministic.
- Persisted conversation memory in a managed Amazon Bedrock AgentCore memory resource rather than process memory, so context survives redeloys and horizontal scaling.
- Instrumented end-to-end observability with LangSmith — tracing every agent run's prompts, tool calls, and replies — and Amazon CloudWatch, tracing the deterministic booking path as well so debugging and testing rely on real execution traces rather than log-guessing.
- Modeled the domain end-to-end and shipped the system as three independently deployed services — Next.js 15 + React 19 frontend on Vercel, FastAPI + SQLite backend on Render, and a LangChain agent on Amazon Bedrock AgentCore — wired together with tightly scoped IAM access.

Tech: Next.js 15, React 19, Tailwind CSS, shadcn/ui, FastAPI, SQLite, LangChain, OpenRouter, Amazon Bedrock AgentCore, LangSmith, Amazon CloudWatch

PROFESSIONAL EXPERIENCE

Senior Quality Engineer (AI-Driven Deployment) — Circle

Remote | Jan 2025 – Apr 2026

- Deployed directly into Circle's FinTech / stablecoin engineering org to architect an AI-driven quality platform using Playwright, Claude Code MCP, and Buildkite MCP, powering autonomous, self-healing end-to-end test systems in production.
- Designed and shipped Selective Test Execution: a Claude Agent SDK-powered system that reads a PR's git diff, reasons about blast radius, and selects only impacted tests on private branches, with full-suite runs on main as a safety net — cutting CI cost by 50% and speeding up developer feedback loops.
- Built intelligent flaky-test recovery and adaptive test-selection strategies that measurably reduced flake rate and CI noise across production suites.
- Engineered automated root-cause analysis that posts actionable fix recommendations directly into pull requests, reducing time-to-resolution for engineering teams.
- Led org-wide adoption of LLM-powered tooling: authored best practices, playbooks, and enablement programs — the same evangelism and enablement work core to FDE and applied-AI roles.

Tech: Playwright, TypeScript, JavaScript, Claude Code MCP, Claude Agent SDK, Buildkite MCP, CI/CD

Engineering Manager, Quality — GitLab Inc.

Remote | 2019 – 2024

- Led a distributed team of 8+ engineers, architecting CI/CD test infrastructure handling 100K+ executions daily and driving pipeline flakiness from a 15% baseline to a stable low-single-digit range.
- Partnered with product and engineering stakeholders to translate ambiguous business requirements into testing strategy — a discovery-to-delivery workflow analogous to FDE client engagements.

Lead SDET / Senior SDET — Freshworks (Freshdesk)

Chennai | 2016 – 2019

- Architected test frameworks spanning 20+ microservices and rolled out Consumer-Driven Contract testing org-wide, preventing breaking changes across service boundaries — systems-integration work core to FDE deployments.
- Delivered SDK test frameworks enabling external developers to build reliable integrations on the Freshworks platform, and represented the org at Rootconf, Global Testing Retreat, and VodQA.

Tech: Selenium, Cucumber, Capybara, Pact, Docker, Jenkins, Ruby, Java, Python

Staff Quality Engineer — PayPal

Chennai | 2015 – 2016

- Enhanced end-to-end frameworks for a global payment-processing platform and codified quality gates and deployment best practices across payment services with cross-functional partners.

Senior Software Design Engineer in Test — Chronus

Chennai | Oct 2013 – Jun 2015

- Designed a web and mobile (Protractor/AngularJS) test automation framework, integrating with Solano CI for rapid feedback loops.

SDET / Application Engineer — Amazon

Chennai | 2009 – 2013

- Reduced Kindle regression execution from 8 hours to 40 minutes (95%) through intelligent parallelization across 100+ test agents — recognized with multiple awards for technical impact.

Member Technical Staff — Juniper Networks

Bangalore | 2007 – 2009

- Developed and tested Hostchecker for Juniper's SSL VPN solutions and resolved critical customer escalations through rapid root-cause debugging.

KEY ACHIEVEMENTS

- AI Innovation: Pioneered self-healing test systems using Playwright, Claude Code MCP, and the Claude Agent SDK at Circle.
- CI Cost Optimization: Shipped an agentic Selective Test Execution system that reads PR diffs and reasons about impact, cutting CI cost by 50%.
- Performance at Scale: Cut Kindle regression execution by 95% (8 hours → 40 minutes) at Amazon via intelligent parallelization.
- Reliability: Drove GitLab CI flakiness from a 15% baseline to a stable low-single-digit range.
- Community Impact: Directed WomenWhoCode Chennai (2018–2024), empowering 2,000+ women in technology through 50+ technical workshops.

SPEAKING ENGAGEMENTS

- DevTestOps: Quality Engineering in Modern DevOps — FOSSASIA Summit, 2023
- DevTestOps: Testing Efficiently in DevOps Culture — The Test Tribe Chennai Meetup, 2022
- Quality Engineering @ GitLab — GDG Sri Lanka, 2019
- Consumer-Driven Contract Testing at Scale — Global Testing Retreat, 2017

EDUCATION & CERTIFICATIONS

- Postgraduate Program in Forward Deployed Engineering & Applied AI Solutions — FDE Academy (Futureense), in progress, expected 2026
- Certification: "AI Evals for Everyone" — 2026
- Director, Women Who Code — Chennai Chapter (2018–2024)