

# Mayank Agrawal

Principal Architect | AI Systems Engineer | Engineering Lead

Bengaluru, India | [me@mayankagrawal.co.in](mailto:me@mayankagrawal.co.in) | +91 6392107195 | [linkedin.com/in/er-mayank](https://www.linkedin.com/in/er-mayank) | [github.com/er-mayanka](https://github.com/er-mayanka) | [Portfolio](#)

**Principal Architect and AI Systems Engineer** with 10+ years building high-throughput distributed systems, leading engineering squads, and shipping AI-native products. Founder of **Zethra OS** (AI-native mobile OS with self-healing LLM patch engine) and **cursor-lite** (fully local multi-model AI coding agent, no cloud APIs). Deep expertise in Node.js, Kafka, AWS, distributed systems, and agentic AI. Scaled platforms to 2M+ users, managed large-scale cloud budgets, and led squads of up to 6 engineers. Open to Staff Engineer, AI Systems Lead, and Engineering Lead roles globally.

## EXPERIENCE

**Zethra OS** | Founder & Lead AI Systems Engineer Mar 2026 - Present · Remote

AI systems company building intelligent developer tooling on local LLMs, with no cloud dependency, no API costs.

- Built **cursor-lite**: a fully local, multi-model AI coding agent powered by Qwen2.5-Coder, Qwen3.5, and LLaMA 3.2, orchestrated via a custom Node.js agent with planning, validation, and step-by-step execution with **zero cloud API** dependency.
- Designed a **stateful tool safety system** that detects missing dependencies, requests user confirmation, installs them, and resumes agent execution without state loss, production-grade reliability from day one.
- Shipped a streaming UI (ChatGPT/Claude-style) with live markdown rendering, syntax-highlighted code blocks, and typewriter reveal built end-to-end in Node.js.
- Architecting **Zethra OS**: an AI-native mobile OS built on Rust and the Linux 6.x kernel. Core feature, **ZethraAI**, is an autonomous system that monitors system health, analyzes crash data, and generates cryptographically signed patches without human intervention.

**Independent Software Consultant** | Software Architecture & Engineering Apr 2024 - Feb 2026 · Remote · US, UK, UAE

- Provided software architecture consulting and technical delivery for enterprise clients across North America, the UK, and the UAE. Engagements are under NDA.
- Developed deep expertise in agentic AI, local LLM orchestration, distributed backend design, and cloud-native infrastructure, directly informing Zethra OS and Cursor-lite.

**Onsurity Technologies** | Team Lead - SDE III May 2022 - Apr 2024 · Bengaluru, KA

B2B healthcare benefits platform providing group health coverage to 2M+ users across India's SME and startup ecosystem.

- Led a 5-engineer cross-functional squad, improving team efficiency by **20%** through structured code reviews and sprint rituals.
- Integrated 3 enterprise HRIS APIs (Keka, Darwinbox, Zoho People), automating employee data sync for enterprise clients and reducing manual onboarding effort by **35%**.
- Architected async job processing pipelines using Kafka, Redis, and BullMQ across 6 microservices (Gateway, RBAC, Identity, Integration, Claims, Notifications), cutting manual processing time by **50%**.
- Scaled AWS and Azure infrastructure from 1M to 2M+ active users while maintaining high availability and cost efficiency within a high six-figure cloud budget.

**Ace Turtle Services** | Senior Software Engineer Sep 2020 - May 2022 · Bengaluru, KA

SaaS B2B omnichannel retail platform (RUBICON) serving Puma, Fossil, Samsonite, and Chicco. Also operated Lee & Wrangler across Flipkart, Amazon, and Myntra.

- Built an asynchronous, event-driven AWS SQS integration that processed millions of marketplace orders during Big Billion Days and End of Season Sale, achieving **zero data loss** under peak traffic..
- Optimised 15+ MongoDB queries across order management and inventory sync flows, reducing API response times and improving platform stability by **40%**.
- Developed Lee & Wrangler merchandising feature bridging offline retail to online channels. Integrated QZ-tray for warehouse invoice printing, achieving **70%** reduction in fulfilment time.
- Led 6 engineers, improving sprint velocity by **30%** and unblocking 12 critical production issues.

**Durvah IT Consulting** | Software Engineer Sep 2019 - Aug 2020 · Bengaluru, KA

- Sole engineer on AirAsia AMS, a live commission invoice automation platform for AirAsia India's finance team and agent network across India. **95%** client satisfaction.
- Migrated database from MongoDB Atlas to on-premises IIS replica set to meet AirAsia's enterprise data security requirements. Implemented Redis caching to maintain performance parity post-migration.

**Promatics Technologies** | Associate Software Developer Mar 2018 - Apr 2019 · Ludhiana, PB

- Sole backend engineer across four production applications: LeQuiz (real-money trivia, UK), ShowUoff (youth sports social platform, UAE), MAC - Meet & Chill (geo-based dating app with real-time Socket.io chat), and TLI Travel (7-language app with PayPal integration); owned end-to-end AWS EC2 infrastructure and DevOps across all systems.
- Took full solo ownership of LeQuiz after team lead departure, **rebuilt Angular frontend from scratch** while maintaining live backend and real-money prize payments.

- Built Node.js REST APIs for the Customer Plan Prediction CRM module, integrating 2 data sources and improving data retrieval speed by 20%.

## KEY PROJECTS

**Zethra OS: AI-Native Mobile OS** - Self / Founder Rust · Linux 6.x · LLMs · Apache 2.0

Next-gen mobile OS with ZethraAI: an autonomous immune system that monitors system health, analyses crash data, and auto-generates cryptographically signed patches. Rust-first userspace for memory safety. Privacy-by-design architecture.

**cursor-lite: Local AI Coding Agent** - Self / Zethra OS Node.js · NestJS · ChromaDB · Qwen2.5-Coder · Qwen3.5 · LLaMA3.2

Fully local multi-model AI coding assistant. Multi-model orchestration layer, stateful tool safety system, and ChatGPT-style streaming UI. No cloud APIs, runs entirely on local LLMs with production-grade Git workflow.

**Onsurity Platform (Gateway, RBAC, Identity, Notifications, LinkedIn scraper)** - Onsurity Technologies Node.js · Kafka · Redis · Playwright · React · Spring Boot · AWS · Azure

Core platform services for a healthcare benefits platform scaling from 1M to 2M+ users. Async Kafka pipelines, enterprise HRIS integrations, and cloud infra across AWS and Azure.

**RUBICON 2.0: Omnichannel OMS** - Ace Turtle Services Node.js · Angular · React · AWS (EC2, SQS, S3, Lambda, DynamoDB) · MongoDB

Omnichannel retail platform serving Puma, Fossil, Samsonite. Async SQS integration processing millions of orders during BBD/EORS. Lee & Wrangler marketplace expansion. QZ-tray invoice printing (70% time saved).

**AirAsia AMS (Agent Management System)** - Durvah IT Consulting Angular 8 · Node.js · MongoDB on-prem · Redis

Live commission invoice automation for AirAsia India's finance team and agent network. On-prem MongoDB replica set migration to meet enterprise security compliance. Delivered solo end-to-end.

**LeQuiz · ShowUoff · MAC · TLI Travel** - Promatics Technologies Node.js · Socket.io · Angular · MySQL · AWS EC2

Four live production apps across real-money trivia gaming, youth sports social, geo-social dating, and multilingual travel safety. Sole backend engineer and DevOps owner across all four. Full architecture ownership on each.

## TECHNICAL SKILLS

<b>AI &amp; LLM</b>	Local LLM orchestration · Multi-model agentic AI · Qwen2.5-Coder · Qwen3.5 · LLaMA3.2 · RAG systems · Prompt engineering
<b>Backend</b>	Node.js (8 yrs) · Fastify · ExpressJS · Java Spring Boot · Kafka · BullMQ · Redis · Socket.io · REST APIs · Microservices
<b>Frontend</b>	React.js · Angular (6 yrs) · TypeScript · Redux · TailwindCSS · Shadcn · Vite
<b>Cloud &amp; DevOps</b>	AWS (7 yrs: SQS, Lambda, EC2, ECS, Beanstalk, IAM) · Azure · Docker · Kubernetes · Ansible · GitHub Actions · CI/CD
<b>Databases</b>	MongoDB (6 yrs) · MySQL · PostgreSQL · Redis · DynamoDB · Sequelize · Mongoose
<b>Observability</b>	Prometheus · Grafana · Loki · New Relic · ELK Stack · Cloudwatch · HotJar
<b>Systems</b>	Rust · Linux 6.x kernel · Distributed systems · System design · Self-healing architecture
<b>Leadership</b>	Technical strategy · Sprint planning · Code review culture · Performance cycles · Cross-functional squad leadership

## EDUCATION & CERTIFICATIONS

**B.Tech, Computer Science & Engineering** - Dr. A.P.J. Abdul Kalam Technical University, Lucknow 2012 - 2016

**Kubernetes Certified Professional** - LinkedIn Learning 2024

**Research Paper** - "Green Technology: A Modern Approach to Make the Environment Hygienic for Human Living." Abstract accepted, CSI Lucknow Chapter National Conference 2015.