

# Vaibhav Pathak

## Senior Backend Engineer

vaibhavpathak9984@gmail.com | +91 9170439308 | Bengaluru, India

[GitHub](#) | [LinkedIn](#) | [Portfolio](#)

---

## Professional Summary

Senior Backend Engineer with 5+ years building distributed systems, microservices, and cloud-native platforms. Focused on AI engineering – RAG pipelines, LLM integration, semantic search, and generative AI in production. Experience leading small teams and shipping products at scale (100k+ MAUs). I care more about what ships than what looks good on paper.

---

## Technical Skills

Languages: JavaScript, TypeScript, Go, Rust, Python

Backend: Node.js, Bun, Hono, Express.js, FastAPI, Django, Flask, gRPC, GraphQL, REST API

Frontend & Mobile: React, Next.js, Flutter, React Native

Databases: PostgreSQL, MySQL, MongoDB, Redis, Amazon Aurora, Memcached

AI / ML: RAG, LLM integration, vector embeddings, NER, generative AI, on-device ML, semantic search, machine learning pipelines

Search: MeiliSearch, Elasticsearch, hybrid semantic + lexical search

Cloud & DevOps: AWS (Lambda, SNS, SQS), GCP, Docker, GitHub Actions, GitLab CI, CI/CD

Messaging: Kafka, Amazon SQS

Tools: Postman, Newman, Xcode, Gradle

---

## Professional Experience

### Senior Software Engineer | Ant Mascot, Bengaluru

Sep 2020 – Present

- Leading all engineering for a \$1.8M B2B company with a 2-person team.
- Architected and built the ERP system from scratch using Python, Node.js, and Go microservices.
- Migrated the entire B2B sales operation to antmascot.com in under 4 months.
- Scaled the platform to 100k monthly active users.
- Launched Sampura Bazaar, a quick-commerce B2C marketplace.
- Engineered the hybrid search and auto-tagging system driving product discovery on the platform.

### Software Engineer | Diago Developers, India

May 2017 – Jul 2020

- Founded and scaled a software services company from 1 to 8 engineers, serving 10+ clients.
  - Collaborated with Parag Arora (Entrepreneur.com 35 Under 35 in AI) on AI-driven products.
  - Built machine learning pipelines for large-scale data processing and auto-scaling backend infrastructure.
  - Developed offline-first Flutter mobile apps with background sync.
- 

## Projects

### Oxyzen – Employee Engagement SaaS

Stack: Go, PostgreSQL, MeiliSearch, AWS | [Website](#) | [Android](#) | [iOS](#)

Private internal social platform for companies – employees share updates, feedback, and work documentation within a closed network.

- Architected multi-tenant infrastructure with database-level sharding per tenant.
- Built a RAG-based AI chatbot trained on company-specific data; updates its knowledge base automatically as new content is added.
- Integrated generative AI for producing company-branded post images.
- Developed on-device autocomplete trained on each user's writing history – no data sent to server.
- Built notification pipeline in Go using AWS SNS + SQS with per-user timezone handling.
- Wrote a Rust CLI for internal dev tooling: code scanning, deployments, and migrations.
- Shipped on Android and iOS.

## Search Infra – Hybrid Search Engine

Stack: MeiliSearch, Python, PostgreSQL

- Replaced Elasticsearch with MeiliSearch for hybrid semantic + lexical search; queries under 50ms.
- Built auto-tagging pipeline on product upload – generates tags, detects overlaps, merges related ones automatically.
- Implemented automatic SEO generation for product tags.
- Integrated NER models to parse actual search intent beyond the raw query.
- Added vector embeddings for product recommendations and similarity matching.
- Deployed full version on Ant Mascot; scaled-down version running on Oxyzen.

## Medzimi – Medical Platform

Stack: Bun, Hono, PostgreSQL, MeiliSearch | [Android](#)

Multi-role medical app connecting patients, doctors, pharmacies, labs, and medical reps.

- Built a handwritten prescription scanner achieving ~94.5% extraction accuracy on real doctor handwriting.
- Developed live medicine search connected to real-time pharmacy inventory.
- Implemented a clinical-parameter-based dose calculator for paediatric patients.
- Built AI-generated personalised diet and workout plans based on user health data, region, and calorie targets.
- Launched in-app community feed across all user roles.

## Fikrabort – Document Extraction & Verification

Stack: Next.js, Go, Redis, PostgreSQL | [Website](#)

Document verification platform for logistics and compliance workflows.

- Built a configurable flow system – define which documents to upload, what data to extract, and what fields need to match across documents.
- Developed a verdict engine that automatically identifies corresponding fields across all uploaded documents and runs matching without manual configuration.
- Implemented cross-document validation for logistics: invoices, airway bills, and permits checked against the same shipment.
- Integrated live VAT and GST validation against government APIs with a clear pass/fail verdict output.
- Built a real-time translation layer for cross-language document matching – Arabic and English documents processed in the same flow.
- Fine-tuned a custom OCR model for mixed Arabic-English documents; off-the-shelf models lacked the accuracy needed.
- Uses multiple LLMs in the extraction pipeline, selected per document type and language.

---

## Achievements

- Shipped Google Assistant integrations in production shortly after the platform launched in 2016.
- Awarded ₹3.8 lakh in cloud credits for early cloud adoption.