

Amal Shaji

Bengaluru, India | amalshajid@gmail.com | amal.sh | linkedin.com/in/amalsh | github.com/amalshaji

OVERVIEW

Full-stack engineer with 5+ years taking systems from zero to production across product interfaces, backend services, data platforms, and AI/LLM infrastructure. Enjoys owning every layer of the system, from React applications and Node.js/Bun services to cloud deployment and observability. Experienced with AWS, GCP, Azure, Cloudflare, and other cloud platforms. Works primarily in TypeScript, Python, Go, and Rust.

EXPERIENCE

Software Engineer II — Plivo

2021 – Present

CPaaS / customer-experience platform

Bengaluru, India

- Owned production features end to end across frontend and backend, building production frontends alongside backend services for customer-data, app-integrations, and voice-agent products.
- Designed an internal customer data platform ingesting real-time website SDK event streams, CSV batch uploads, and Salesforce/Shopify syncs; performed identity (entity) resolution across sources into a unified customer profile using Python, Jitsu, and ClickHouse.
- Rewrote the voice-agent platform's LLM service from Python to TypeScript on Bun, reducing service-side response latency by 30% and adding in-flight query cancellation to handle speech interruptions in real time.
- Worked with external partners to migrate LLM workloads from OpenAI to Azure OpenAI, enabling the company to use \$250K in startup credits.
- Optimized the CI/CD pipeline for an internal service, cutting build times by 70-80% - from 10 minutes to 2-3 minutes - and reducing CI infrastructure costs.
- Standardized OpenTelemetry instrumentation across all microservices, enabling end-to-end distributed tracing across service boundaries to diagnose failures and latency bottlenecks.
- Built an end-to-end integrations platform spanning 30+ third-party providers, including production frontend workflows, OAuth2/API-key authentication, provider-specific MCP tools, and document syncing into RAG pipelines for LLM-powered features.
- Created agent-observability, an open-source verification layer for voice agents: session tracing, simulation-based evals, and LLM-as-judge conversation judgments used to validate agent behaviour before and after release.
- Fine-tuned Qwen-family models end-to-end for internal use cases, covering dataset preparation, training runs, and evaluation, achieving 90-94% accuracy on target scenarios.

Technical Author, TestDriven.io

Dec 2021 – Mar 2024

- Authored 16 hands-on tutorials on production Python web engineering, covering authentication and security, FastAPI/Flask/Django, Dockerized deployments, testing and concurrency, and frontend and ML integrations. testdriven.io/authors/shaji/

Research Intern — Leading India, Bennett University

Jun 2020 – Jul 2020

- Studied node-deployment strategies for heterogeneous wireless sensor networks to improve network lifetime.

OPEN-SOURCE PROJECTS

- portr** (Go, TypeScript) — self-hosted tunneling tool exposing local HTTP/TCP/WebSocket services publicly; connection multiplexing over a single control channel. 3k+ GitHub stars. github.com/amalshaji/portr
- DBCoper** (Rust, Tauri, React) — desktop database client for PostgreSQL, SQLite, Redis, and ClickHouse. 400+ GitHub stars. github.com/amalshaji/dbcooper

TECHNICAL SKILLS

Languages: Python, TypeScript, Go, Rust, SQL

Data & Streaming: ClickHouse, PostgreSQL, Redis (caching, queues), Jitsu event pipelines, SQLite

AI / LLM Systems: MCP, agent evals & observability, LLM fine-tuning (Qwen), LiveKit Agents, Pipecat, SIP/RTP

Cloud & Infrastructure: AWS, GCP, Azure, Cloudflare, Docker

Frameworks & Runtimes: Node.js, Bun, React, FastAPI, Tauri, Streamlit

EDUCATION

Indian Institute of Information Technology, Gwalior — B.Tech, Computer Science

2017 – 2021