

Ayush Jaipuriyar

+91 9968302859 | ayushjaipuriyar21@gmail.com | Bangalore, India
portfolio.ayushjaipuriyar.uk | linkedin.com/in/ayushjaipuriyar | github.com/ayushjaipuriyar

PROFESSIONAL SUMMARY

Full-stack software engineer with 2+ years of experience designing and shipping distributed systems, microservices architectures, and cloud-native web and mobile applications. Proficient in TypeScript, Python, Node.js, NestJS, React, Next.js, React Native, PostgreSQL, MongoDB, Redis, Kafka, Docker, Kubernetes, AWS, Cloudflare, and CI/CD with GitHub Actions. Experienced in system design, API development (REST, GraphQL, WebSockets), performance optimization, and Agile delivery.

EDUCATION

University of Glasgow
M.Sc. in Computer Science

Sep. 2024 – Dec. 2025
Glasgow, UK

EXPERIENCE

Full Stack Software Developer
Seabuddy

Feb. 2026 – Present
Singapore

- Architected a RAG-based AI assistant using Cloudflare AutoRAG, Workers, and Hono, enabling ~300 DAU to query policy documents instantly, cutting average manual lookup time by ~70%.
- Developed a coin-based gamification system (tasks, events, streaks) that grew DAU by ~20% (from ~250 to ~300 active users) within 60 days of launch.
- Designed a multi-tenant platform architecture with per-tenant isolated databases and async PDF ingestion pipelines (CF Queues, R2, KV), enabling onboarding of multiple shipping companies with zero cross-tenant data leakage.
- Optimized backend and database query plans, reducing P50 API latency from 320ms to 140ms (56% improvement) and improving throughput under peak load.
- Shipped a React Native app (iOS/Android) with social feed, events, and real-time group chat via WebSockets, serving 1,000+ users with support for 200+ concurrent connections.

Full Stack Software Engineer
Healthtrip

Jan. 2024 – Sep. 2024
Noida, India

- Increased API throughput from 120 to 210 RPS by migrating a PHP monolith to NestJS microservices, improving system reliability by 75%.
- Created REST, GraphQL, and WebSocket APIs serving 8,000+ DAU, reducing P50 response times by 45% (220ms → 120ms) through query optimization and connection pooling.
- Implemented Redis caching layer, reducing database reads by 70% and cutting read latency from 300ms to 30ms (90% improvement).
- Introduced a custom i18n translation system covering 9 languages, driving 75% engagement growth (5K → 8.75K DAU) across non-English markets.
- Spearheaded multilingual Elasticsearch integration, accelerating search throughput from 50 to 80 results/sec (60% improvement).
- Deployed a Kafka-based notification pipeline processing 10,000+ events/day, cutting P95 delivery latency by 30%.
- Automated CI/CD pipelines using GitHub Actions and Docker across AWS and Cloudflare, reducing deployment cycle from ~2 hours to under 10 minutes (~90%) and cutting page load time by 50% (3.2s → 1.6s).

Software Developer
AST Consulting

Jun. 2023 – Dec. 2023
New Delhi, India

- Launched a GenAI SaaS platform for social media content generation (React, NestJS, MongoDB), onboarding 200+ users within 3 months and supporting 50+ concurrent content generation requests.
- Built REST and GraphQL APIs in NestJS integrating OpenAI GPT models, with retry logic and token-budget management, achieving <500ms median response time under production load.
- Integrated Stripe and Chargebee subscription billing, adopted by 25% of users within launch quarter, and automated Telegram bot notification workflows.
- Optimized MongoDB aggregation queries and caching layers, reducing latency from 320ms to 180ms (44%) for data-heavy analytics dashboards.
- Provisioned AWS infrastructure with GitHub Actions CI/CD and CloudFront CDN, reducing image load times by 50%, accelerating deployments by 40%, and sustaining 95%+ uptime.

PROJECTS

- LeetCode MCP Server** | *Node.js, TypeScript, Anthropic, OpenAI* GitHub
- Built and published an npm MCP server exposing LeetCode GraphQL endpoints for problems, user profiles, and code submission; implemented submission relay with polling, caching, and rate-limit handling.
- TurboM3U8** | *Python, ffmpeg, tkinter, PyQt* GitHub
- Built a multi-threaded CLI tool for downloading and assembling HLS .m3u8 streaming segments with robust error handling and recovery logic, improving download reliability by 70% vs single-threaded baseline.
- Near-RT RIC ML-Based Malicious Traffic Detection** | *Python, PyTorch, TensorFlow, O-RAN, ML* GitHub
- Developed an ML pipeline for real-time O-RAN traffic analysis with live ingestion and monitoring, achieving **67–73%** accuracy and up to **76%** F1-score on malicious traffic classification.
- Lenovo Vantage-Linux: ACPI Tuning Utility** | *Linux, Unix, ACPI, Bash* GitHub
- Devised a command-line Linux utility exposing ACPI performance tuning controls, enabling thermal profile switching and improving battery life by **20–30%**.

TECHNICAL SKILLS

Languages: TypeScript, JavaScript, Python, Java, SQL, Go, C/C++, Bash, PHP

Frontend: React, Next.js, Redux, React Native, Expo, HTML, CSS, Responsive Design

Backend: Node.js, NestJS, Express, Hono, Django, Flask, REST APIs, GraphQL, WebSockets, Microservices, API Design

Databases: PostgreSQL, MySQL, MongoDB, Redis, Firebase, Elasticsearch, BigQuery, Query Optimization

Cloud & DevOps: AWS (EC2, Lambda, S3, CloudFront, RDS), GCP, Cloudflare (Workers, R2, KV, Queues, AutoRAG), Docker, Kubernetes, Terraform, GitHub Actions, Jenkins, CI/CD, IaC

Messaging & Caching: Kafka, Redis, BullMQ, Cloudflare Queues, Pub/Sub

AI/ML: OpenAI API, Anthropic API, RAG, PyTorch, TensorFlow, LLM Integration, AI/ML Integration

Testing: Jest, PyTest, JUnit, Supertest, Postman, Unit Testing, Integration Testing

Tools: Git, Sequelize, Jira, Linux, Bash, Agile/Scrum, Technical Documentation