

# Kuldeep Singh

+91-8437077071 | Email: [kuldeep8437.singh@gmail.com](mailto:kuldeep8437.singh@gmail.com) | [linkedin/kuldeepsingh](https://www.linkedin.com/in/kuldeepsingh) | [github/kuldeepsingh](https://github.com/kuldeepsingh)

## TECHNICAL SKILLS

---

**Languages:** Java, Python, C/C++, JavaScript, Go

**Tools & Technologies:** Git, Docker, Kubernetes, CI/CD, AWS, Redis, MongoDB (NoSQL), PostgreSQL (SQL), GCP

**AI/ML:** LangChain, LLMs, Ollama, MCP (Model Context Protocol), AWS Bedrock, Gemini

**CourseWork/Skills:** Data Structures and Algorithms, Operating Systems, Computer Networks, OOPS Concepts, Web Development, LLD (Low level design), Clean Architecture, HLD (High level design)

## WORK EXPERIENCE

---

### Software Development Engineer

09/2023 - Present

*Samsung R&D Delhi*

*Noida, India*

- Migrated high-scale core backend microservices (100M+ daily API requests) from Node.js to Java Spring Boot, containerized from EC2 to EKS, and implemented GitHub CI/CD — boosting throughput/stability and cutting infra costs by **\$18000** per month
- Redesigned observability stack from ELK to S3 + Athena + Superset — reduced logging/monitoring costs by **\$15000** per month while preserving real-time dashboards.
- Led development of an internal LLM portal for 100+ service owners to manage AWS, GCP, Azure via natural language — cut manual console effort by **70%**.
- Delivered plug-and-play MCP servers for internal teams to query AWS Trusted Advisor and Cost Explorer via natural language, drove **\$300,000+** in cloud cost savings.
- Spearheaded Redis-based asynchronous job processing service (5M+ jobs/day) with multi-threading in Spring Boot; integrated Samsung's DKMS encryption to secure data of 10M+ users.
- Mentored two engineers, reviewed designs/code, and drove architecture decisions for critical backend components.

### Software Engineer Intern

01/2023 - 06/2023

*Vitto Microfinance AI*

*Gurgaon, India*

- Co-architected Notification Management microservice in TypeScript — processed **500K+** notifications/month with 99.9% delivery success.
- Implemented Redis caching with rate limiting and targeted invalidation strategies, achieved **90%** faster repeated reads, improving query performance.
- Developed a high-performance QR payment feature for UPI payments with YES BANK as the payment aggregator (10,000+ transactions/month, **INR 10M** value).
- Optimized cold-start and deploy speed by moving from Ubuntu to Alpine slimming Docker images by **65–70%**.

## PROJECTS

---

### Financial AI Assistant | Python, MCP, LLMs (Gemini), React, Docker

- Built a full-stack financial assistant that integrates a custom Splitwise MCP server, FastAPI backend, and React chat UI to let users manage and analyze expenses via natural language.
- Implemented an MCP server Wrapped 20+ Splitwise APIs with OAuth and abstraction layer for LLM integration; optimized token usage by 70% via smart context management.

## EDUCATION

---

### JC Bose University of Science and Technology YMCA

Faridabad, India

*Bachelor of Technology (ECE)*

*08/2019 - 06/2023*

## ACHIEVEMENTS

---

- Delivered cumulative cloud infrastructure cost savings of up to **\$40,000** per month across AWS cloud services.
- Solved 800+ DSA problems across different platforms. Qualified Samsung's SWC Professional on the first attempt with **98+** percentile.
- Achieved 5 Stars on CodeChef and global rank 26 in Starters-61. Also secured AIR 320 in Google Kickstart 2022.