

VARUN CHOUDHARY

Software Development Engineer

chdhryvarun@gmail.com | +352 691323196 | Luxembourg, Luxembourg | [linkedin.com/in/varun-choudhary-baa455b6](https://www.linkedin.com/in/varun-choudhary-baa455b6)

SDE with **8+ years** architecting distributed systems at scale. Currently building **agentic AI platforms** at Amazon — multi-agent LLM systems, custom MCP servers, and RAG pipelines powering the EU Quick Commerce org.

WORK EXPERIENCE

Software Development Engineer II, *Amazon*

11/2020 – Present | Luxembourg & Gurgaon

- **QC Expansion Assistant — Multi-Agent System (2026):** Led the design of a multi-agent automation platform to accelerate Quick Commerce node launches across EU marketplaces. Authored the architecture doc, selected the agent framework (**Strands Agents SDK + GraphBuilder DAG + Bedrock Claude**), and defined the two-phase auth model (client-side impersonation → remote delegation). Designed dependency DAG with steering-hook guardrails, rollback strategy, and Evals SDK test cases. **Projected impact: reduce per-node SDE effort from 3–4 weeks to <3 days.** Mentored SDE team on implementation.
- **Agentic AI Platform for EU Quick Commerce (2026):** Architected and shipped a production agentic AI platform serving the hyperlocal delivery org. Designed a custom **MCP server exposing 18 enterprise tools** and 8 agentic business flows, plus a RAG-grounded Knowledge Agent that eliminated LLM hallucinations. Code reviewed and approved by principal engineers and AI specialists.
- **Generative AI Config Migration:** Led migration of hard-coded configurations across **10+ mission-critical services** to a dynamic config store using **GenAI-assisted code generation**, reducing feature delivery time by **80%** and eliminating deploy-to-toggle dependencies.
- **Seller Central UX Revamp + ACS Migration (2025):** Architected React-based UX revamp for merchant-fulfilled orders across UK, DE, JP, US (**+30% seller engagement**); led **Amazon Central Services** platform migration across JP/UK/DE via dual-write progressive dialup — zero-downtime across millions of seller orders.
- **Transportation Compliance Platform — India GST E-Way Bill (2022–2024):** Architected event-driven compliance system for India's legally mandated nation-scale regulatory platform. Designed idempotent pipeline (SQS + DynamoDB + Lambda + ECS) integrating with government APIs for real-time filing. Delivered Phase 2–5 rollouts covering Last-Mile, Rail, Air-Freight Transshipment. Built aggregator service with S3 data lake, Glue ETL, Redshift Spectrum. **Zero high-severity incidents over 18 months**; processes millions of shipments.
- **Secure Rejects with OTP (Jan–Apr 2024):** Led cross-team package-delivery security initiative — scoping, HLD, event-routing integration — achieving a **75% reduction in delivery security incidents**. Extensible design now supports AMZL OTP, Ship with Amazon, and global re-trigger workflows.
- **Production Reliability & Leadership:** Drove reliability initiatives across **40+ services**; deprecated a 10-year-old tier-1 legacy service; resolved multiple P1/P2 incidents via deep-dive RCAs. Mentored 3 interns (**100% SDE conversion**), supported 6 SDEs (**4 promoted, 1 in progress**), conducted **50+ interviews** including bar-raiser calibration.

Software Development Engineer, *Limeroad.com*

06/2018 – 10/2020 | Gurgaon, India

Built the complete backend for thrift and brand store platforms, with highly configurable listing pages powered by **Solr**. Revamped the user login flow to remove mandatory authentication, reducing drop-off by **17%**. Integrated a **feeds-based Facebook paid channel system** and optimized a **dynamic pricing engine** — resulting in an **11% GMV increase**.

PERSONAL PROJECTS

NSE Multi-Agent Stock Analyzer, github.com/varunch/nse_agent

2026

Production-grade hybrid AI trading system for NSE India and S&P 500. End-to-end agentic AI design with safety-critical execution. Stack: Python, CrewAI, Strands, MCP, ChromaDB, Bedrock / OpenAI / Ollama.

- **Architecture:** 5-layer system — deterministic orchestrator, 7 rule-based agents, CrewAI + Strands LLM layer, ChromaDB RAG, 2 custom MCP servers. Hybrid pipeline: rule-based screening → multi-agent deep analysis (Pydantic-validated) → portfolio risk.
- **Safety & Flexibility:** code-level order caps, dedup window, paper-trading mode, emergency stop, read-only live mode. Runtime-switchable Bedrock / OpenAI / Ollama; RAG freshness policy (90d stale, 180d drop).

TECHNICAL SKILLS

AI / ML

LLMs (GPT, Claude, Bedrock), Agentic AI, MCP, RAG, CrewAI, Strands, LangChain, Prompt Engineering, LLM Evals, ChromaDB

Cloud / AWS

Lambda, ECS, Step Functions, S3, SQS, SNS, DynamoDB, CloudWatch, Glue, Redshift, Bedrock, CDK, KMS

Databases

DynamoDB, MySQL, ChromaDB, Redshift

Languages

Java (expert), Python, TypeScript, JavaScript

Web / Frameworks

React, Spring / Spring MVC, Node.js

Practices

Event-driven microservices, Distributed systems, Systems design, CI/CD, IaC, Security-by-design

EDUCATION

B.Tech Computer Engineering, *Delhi Technological University*

08/2014 – 05/2018 | Delhi, India