

Sai Sushanth Thakkalapally

Hyderabad, India | +91 6303212519 | saisushantht@gmail.com | linkedin.com/in/saisushantht | saisushantht.com

Professional Summary

Backend SDE with **4+ years at Amazon** (FinTech, Returns) shipping global compliance and platform systems. Delivered **DAC7/Inform Act** workflows for **670K+ tax identities** across **8 partners**, led **JDK 17 migration** across a monolith with **90+ downstream clients** at zero downtime, and built a **LangGraph/RAG multi-agent AI system** that cut RCA MTTR by **67%**. Two-times **Reliability Award** winner.

Technical Skills

Languages: Java, Go, Python

Backend: Spring Boot, REST APIs, Microservices, Distributed Systems, API Design, Event-Driven Architecture, SQS, SNS

Databases: DynamoDB, MySQL, PostgreSQL, MongoDB, Redis, Query Tuning

Cloud / DevOps: AWS (EC2, S3, Lambda, SQS, SNS, DynamoDB, X-Ray, CloudWatch, IAM, SES, ECS, OpenSearch/Elasticsearch), Docker, AWS CDK, CloudFormation, CI/CD, Infrastructure as Code

Observability: AWS X-Ray, CloudWatch Metrics & Logs, Distributed Tracing, Production Debugging, RCA

Practices: System Design (HLD/LLD), Concurrency, Caching, Operational Excellence, Scalability Engineering, Agile, Code Reviews, Secure Coding

AI / Automation: Multi-Agent Systems, LangGraph, LangChain, RAG Pipelines, MCP Servers, LLM Integration

Experience

Amazon

Hyderabad, India

Software Development Engineer II

Dec 2024 – Mar 2026

- Reduced region onboarding from **weeks to hours** by migrating **25+ non-US compliance workflows** to a **runtime configuration-based architecture** enabling deployment-free launches.
- Delivered zero-downtime **DAC7** and **Inform Act** compliance workflows ahead of regulatory deadlines, serving **670K+ tax identities** across **8 external partners** and averting multi-jurisdiction non-compliance penalties.
- Authored **HLD** for a unified partner API with least-privilege data access, consolidating **5 endpoints into 1** for **30K+ external client calls/month**; design doc cited internally for quality.
- Drove **JDK 8/11 to 17 migration** across **4 services** including a monolith (**3K+ build issues, 90+ downstream clients**) with **zero production downtime**.
- Delivered zero-downtime **cross-region traffic migration** for a high-availability service, closing IAM, auth, and monitoring gaps across **15+ downstream dependencies** via **IaC-only governance and drift detection**.
- Cut RCA MTTR from **45 to 15 minutes (67%)** by building a **LangGraph/RAG multi-agent AI system** correlating codebase context, business rules, and live CloudWatch logs.

Software Development Engineer I

Jul 2022 – Nov 2024

- Delivered SBA-DSP VAT compliance workflow unlocking **£233MM** in customer VAT savings across UK marketplaces.
- Engineered a **data prefill engine** for tax questionnaires, reducing customer completion time by **25%** across **65+ global businesses**.
- Designed **LLD** and delivered **first-of-its-kind non-withholding compliance solution** supporting **100K+ payees** with backward compatibility across **75 existing workflows**.
- Implemented **per-client rate limiting** on core API serving **90+ consumers**, reducing spike-related errors by **90%** and preventing cascading failures.
- Resolved **490+ incidents (31 high-severity)**, eliminated duplicate alerting, mentored **3 engineers/interns**, and drove team code quality through **380+ code reviews**.
- Built Java/DynamoDB/SES notification workflow with GSI-based paginated audit extraction over **100K records**; cut query latency **50%** via indexed access patterns and caching. (*Intern, Feb–Jun 2022*)

Awards & Recognition

Reliability Award – Amazon, two consecutive years (2023, 2024) | **Operational Excellence Champion** – Amazon

Projects

GoKV — LSM-Tree Key-Value Storage Engine (Go)

GitHub

- Pure Go **LSM-tree KV store** with WAL, skip-list memtable, SSTables, Bloom filters, and compaction; benchmarks of **100K+ writes/s** and **200K+ reads/s** on M1.

ContractLens — AI Contract Analyzer

Live | GitHub

- **Multi-agent AI web app** flagging risky clauses and explaining legal jargon via a **4-agent LangGraph pipeline** with RAG, vector DB retrieval, and a React frontend; deployed on Railway + Vercel.

Education

Neil Gogte Institute of Technology

Hyderabad, India

Bachelor of Engineering in Computer Science and Engineering, CGPA: **8.27**

May 2018 – Apr 2022