

Sivasanker Nandakumar Padmapriya

Portfolio: sivasanker.com

GitHub: github.com/DevelopedBy-Siva

Email: its.sivasanker@gmail.com

Mobile: +1 (623) 270-0065

LinkedIn: linkedin.com/in/sivasanker-np

WORK EXPERIENCE

Senior Software Engineer | *Promoted from Software Engineer in 2022*

Sep 2020 – Dec 2024

Wipro Limited

India

- Built unified shipment-processing microservices (Spring Boot, Kafka) across FedEx and TNT integrations, scaling to 150K+ requests/hour and delivering \$2M+ in annual operational savings.
- Designed UI/UX and built a full-stack device management platform (React, Flask) to centralize configuration, authorization, and administration for 2,600+ workstations, reducing manual IT effort by 80%.
- Redesigned CI/CD pipelines (GitHub Actions, Docker, AKS) for 15 microservices, integrating automated testing, code-quality gates (SonarQube), and release workflows to reduce deployment time by 55%.
- Diagnosed and resolved performance bottlenecks in core logistics APIs by eliminating N+1 query patterns and introducing Caffeine caching, reducing database round trips by 80%+ and improving p95 latency by 25%.
- Engineered resilient carrier integrations using bounded retries, exponential backoff, and circuit breakers, reducing carrier-related shipment failures by 60%.
- Built an Okta-integrated API gateway that centralized JWT authentication and role-based access control (RBAC) across microservices.

TECHNICAL PROJECTS

AIOps Incident Detection & Diagnosis System | LangChain, FastAPI, PostgreSQL, React

[GitHub](#)

- Engineered a full-stack incident detection system that analyzes Grafana Loki logs, groups related failures into incidents, and combines runbooks with LLM reasoning to diagnose root causes.
- Implemented a policy engine to control the LLM's automated actions, achieving 92% diagnostic accuracy across 170 evaluation cases.

GitHub Codebase RAG Assistant | AWS Bedrock, Chroma, Docker, GitHub Actions, React

[GitHub](#)

- Developed and deployed a RAG application for querying GitHub repositories, featuring code-aware indexing and vector storage to deliver source-cited answers.
- Engineered a hybrid retrieval pipeline (BM25 + semantic search + reranking), achieving an 83% top-5 hit rate and 0.92 faithfulness across 3 open-source repositories.

LLM Fine-tuning for Code Generation | PyTorch, LoRA

[GitHub](#)

- Fine-tuned Qwen2.5-Coder-14B (LoRA) for a local coding assistant, raising HumanEval pass@1 from 78% to 86%.
- Curated 84K Python training samples from 5 open-source datasets through filtering, cleaning, and deduplication.

TECHNICAL SKILLS

Languages & Frameworks: Python, Java, TypeScript, JavaScript, SQL, FastAPI, Spring Boot, Node.js, React

AI/ML: PyTorch, RAG, Agentic AI, LLM Evaluation, Fine-tuning, LangChain/LangGraph, MCP

Cloud & Data: AWS (EC2, S3, DynamoDB, Lambda, Bedrock), Azure (AKS), PostgreSQL/pgvector, Redis

Tools & Systems: Docker, Kubernetes, Linux, Kafka, Grafana, GitHub Actions, Git, REST APIs, GraphQL

EDUCATION

M.S. in Computer Science – Arizona State University, GPA: 3.96/4.0

Jan 2025 – Dec 2026

B.Tech. in Computer Science & Engineering – SRM Institute of Science and Technology

2016 – 2020

ADDITIONAL

- **Teaching Assistant, Arizona State University (Aug 2025 – Apr 2026):** Graded exams and projects for 450+ students across Programming Languages, Software Engineering, and Theoretical Computer Science.
- **Certifications:** IBM AI Developer Professional (2026); Meta Front-End Developer Professional (2025)