

VYSHNAVI D P

☎ +1 669 340 9876 ✉ vyshnavi.dyvandinnepullareddy@gmail.com [in linkedin](#) [github](#) San Jose, CA 95126

Summary

Software engineer with 6+ years of experience in distributed systems and backend development. Built scalable backend services using containerized infrastructure (Kubernetes/Docker) and serverless architectures on AWS, and developed agent-based LLM systems for planning, tool orchestration, and multi-step workflow execution.

Professional Experience

San José State University

Sep 2025 – Present

Research Assistant

San Jose, CA

- Built an autonomous agent system in **Python** and **LangChain** with a planner-executor architecture, structured tool registry, and step-level reflection/retry loop, evaluated on multi-step reasoning and web-navigation tasks.
- Designed an evaluation harness comparing **LLaMA-3** and **Mistral** variants on tool-call accuracy, JSON schema adherence, and error-recovery, with per-step traces and a failure-mode taxonomy that informed prompt and tool-interface redesign.

Accenture

Aug 2021 – Jul 2024

Senior Software Engineer

Bangalore, India

- Engineered backend microservices for device diagnostics and telemetry workflows in **Java/Spring Boot** on **AWS**, serving a managed fleet of **10M+ PCs** across diagnostic, scan-orchestration, and warranty-validation flows.
- Re-architected status delivery for long-running diagnostic jobs from **client-side polling** to **WebSocket push streaming** with **Java async processing**, removing steady-state poll traffic and improving tail latency for active scan requests.
- Architected robust event-driven orchestration services for hardware scans and warranty validation using **Spring Boot** and **AWS SQS/SNS** messaging, enabling highly reliable execution across millions of daily health checks.
- Built distributed telemetry processing pipelines to analyze device diagnostic signals at scale, using **Java**, **Spark**, and **Hive** to process data across **10M+ devices** and improve early issue detection.
- Deployed **Dockerized Spring Boot** services on **Amazon EKS** using custom **Helm charts**, **ConfigMaps**, and environment-specific **Kubernetes** settings, standardizing deployments across dev, QA, and production.

Tata Consultancy Services

Jun 2018 – Aug 2021

Software Engineer

Chennai, India

- Developed modular backend services for a retail optimization platform using **Java Spring Boot**, successfully processing **1,000+ business parameters** to generate space-planning recommendations within seconds for **1,400+ stores**.
- Engineered distributed search and optimization APIs incorporating **Redis caching** and concurrent **Java multithreading**, maintaining **sub-200ms P95 latency** across **2M+ SKU records** during high-concurrency planning workloads.
- Implemented multi-tenant SSO and session management with **Spring Security OAuth2**, **JWT**, and **Redis-backed sessions** across customer-facing applications, with token rotation, tenant-scoped claims, and a shared logout channel.
- Automated provisioning of cloud demo environments using **AWS**, **Terraform**, and **Jenkins** pipelines, cutting environment spin-up from **45 min to 18 min** and making per-customer demos reproducible from version-controlled infrastructure.
- Led security and regression testing for a major version migration, adding **JUnit/Mockito** coverage for authentication, tenant isolation, and pricing logic, and remediating critical **OWASP** findings in access control, deserialization, and stored XSS.

Technical Skills

Programming Languages: Java, Python, Go (Golang), C++, JavaScript, TypeScript, SQL

Backend & APIs: Spring Boot, Spring Security OAuth2, FastAPI, RESTful APIs, gRPC, WebSockets, GraphQL

Cloud & Infrastructure: AWS (EC2, S3, RDS, ALB, DynamoDB, SQS, SNS, EKS), Docker, Kubernetes, Helm, Terraform

Databases & Big Data: PostgreSQL, MySQL, Redis, DynamoDB, pgvector, Spark, Hive, Presto, MapReduce

Testing & Observability: JUnit, Mockito, Jenkins, GitHub Actions, Prometheus, Grafana, SonarQube, OWASP ZAP, Git

AI / LLM Systems: LangChain, Pydantic, Hugging Face Transformers, Agents, Tool Calling, Embeddings, Vector Search

Education

San Jose State University

Aug 2024 – Present

Master of Science in Software Engineering (CSE) | Expected May 2026

San Jose, CA

- **Relevant Coursework:** Machine Learning, Deep Learning, Natural Language Processing (NLP), Data Mining, Enterprise Distributed Systems, Enterprise Software Platforms

Projects

Helios | Go, Distributed Systems, LSM Tree, Raft, PromQL

[-link](#)

- Built a **distributed time-series database in Go** with an **LSM-tree storage engine** (WAL, memtable, SSTables, leveled compaction) and **Gorilla compression**, achieving **~10x** compression for write-heavy time-series workloads.
- Architected a **3-node** sharded cluster with **Raft-based replication**, **Jump Consistent Hashing** for **O(1)** shard routing, a **PromQL-compatible** query layer, and in-engine **EWMA anomaly detection** for real-time alerting.

Sentinel | Python, FastAPI, PostgreSQL, pgvector, GitHub App, BM25, RRF

[-link](#)

- Built an **AI-powered code review platform** with a **FastAPI** backend, **PostgreSQL + pgvector** retrieval layer, and **GitHub App** integration, using hybrid **BM25 + dense retrieval** with **RRF fusion** for context-aware pull request reviews.
- Engineered a **reproducible evaluation harness** for **100 hand-labeled pull requests** with **precision/recall/F1** regression gates, enabling prompt iteration, cost-efficient LLM reviews, and CI-based quality checks.