

Sandeep Chavan

Product Manager · Enterprise SaaS & AI Systems · Stakeholder & Vendor Management

📍 Hyderabad, India ✉ sandeep.chavan422@gmail.com 📞 +91 98495 91933

🔗 sandeep-chavan-portfolio.vercel.app 🌐 sandeep-co4 📱 224Sand

Summary

Also built: recruiter-portal-mu.vercel.app, a live recruiter-facing JD/profile alignment tool.

Technical Product Manager with 9 years delivering enterprise software at scale, owned a CPQ platform serving 1,500 sales engineers across global markets at Infosys; directed \$1M-\$3M cross-functional programs for DHL across US logistics operations. Builds independently in AI systems as a side practice, multi-agent pipelines, fault-tolerant LLM infrastructure, and reliability engineering, bringing hands-on technical depth to product and program leadership. Most effective where product vision and technical delivery stay tightly coupled across complex, multi-stakeholder environments.

Experience

Senior Consultant - Product & Engineering, Infosys – Hyderabad, India Mar 2024 - Mar 2026

- Owned end-to-end delivery of a multi-tenant CPQ platform for 1,500 customers across regions, running a 60-item backlog through Agile/Scrum and shipping the MVP in 12 weeks.
- Worked with engineering to bring defect rates down from 3% to under 0.5% and cut data replication latency from 48 hours to 12, mostly by tightening the release-review process and pushing for better telemetry.

Senior Consultant, Infosys (Pepsico account) – Hyderabad, India May 2023 - Mar 2024

- Managed end-to-end deployment of manufacturing applications across 8+ Pepsico plants, coordinating IT, manufacturing, and project management teams to keep rollouts on schedule.
- Owned the QA gate before go-live, validated each deployment against functional requirements and industry standards, cutting post-launch defect tickets by roughly 15% versus earlier rollouts.

Technical Program Manager, DHL (Contract) – Phoenix, AZ, USA Feb 2022 - Jan 2023

- Ran SAFe programs in the \$1-3M range end to end, requirements through post-launch, delivered on time and about 10% under budget, across teams and stakeholders spanning multiple time zones.
- Reworked the partner onboarding process, cutting it from 8 weeks to 3.5 and bringing on 4 partners against an original target of 3.
- Owned vendor relationships and cross-team alignment for a distributed program; churn dropped 8% over the engagement.

Product & Software Manager, Cricket Wireless – Seattle, WA, USA Oct 2020 - Jan 2022

- Owned automation and process improvements across 20+ retail locations, lifting operational efficiency 18% and cutting losses by roughly a third.
- Sustained a 25% KPI improvement over 12 months through tighter program governance and more consistent cross-team accountability.
- Led system migration delivery with full alignment between engineering and business units across every milestone, no rollbacks, no alignment failures.

Business Data Analyst, NRC Health – Atlanta, GA, USA May 2018 - Apr 2020

- Built nine operational dashboards in T-SQL and Tableau that became the team's day-to-day source of truth, lifting service scores 15% and data accuracy 20%.
- Automated reporting that had previously been a manual weekly process, freeing up most of a day each week for the team.
- Cut data retrieval latency 12% through T-SQL optimization and BI architecture work.

Senior Software Engineer, Infosys McCamish – Hyderabad, India Sep 2015 - Dec 2017

- Maintained 99.9% processing accuracy on a regulated platform across three compliance audits with zero findings.
- Cut document processing time 15% through SDLC and pipeline improvements.
- Wrote and maintained production code on a regulated enterprise platform, zero compliance failures across three audit cycles.

Software Engineer, IBM – Hyderabad, India Jun 2015 - Sep 2015

- Reduced ETL errors 20%, work that held up as the foundation for reporting projects in later roles.

Select AI & Systems Projects

AI Infrastructure & Reliability Engineering

Python (open source, MIT-licensed)

- Published a fault-tolerant multi-provider LLM router that fails over across providers, time-boxes rate-limit disables instead of stranding a pipeline on its slowest fallback, and routes in cost order, keeps long-running autonomous agent runs alive on free and low-cost tiers.

Multi-Agent AI Orchestration Platform

Python · LangGraph

- Architected a 19-agent orchestration system on LangGraph, each agent running its own generate-verify-refine loop with a shared QA/QC/Legal governance gate and a 3-retry ceiling before anything escalates instead of silently failing.

AI Trust & Identity Verification Platform

Python · Next.js

- Designed and shipped a two-sided rental marketplace with AI-based identity verification, OCR plus fraud detection clears a government ID in under 60 seconds, replacing physical deposits with a digital trust score, live across multiple cities with a separate vendor portal.

AI Candidate-Role Alignment Assessor

Next.js · Google Gemini API

- Built and shipped a recruiter-facing web tool where any recruiter pastes a live job description and gets an instant, evidence-based Role Alignment Matrix scored against my real profile, one row per distinct JD requirement with no fixed cap.

Automated Media Scoring & Generation Pipeline

Python · Local Vision & Language Models · FFmpeg

- Built a local-first computer-vision scoring pipeline that ranks incoming media across eight weighted quality axes, using an on-device vision model with a deterministic heuristic fallback so scoring never depends on a paid API.

Multi-Model Identity-Consistent Generation Pipeline

Python · Fooocus (SDXL) · FaceFusion · FastAPI · Next.js

- Architected a three-process, three-virtualenv system that isolates a lightweight orchestrator from two heavy ML runtimes with conflicting pinned dependencies, dispatching every generation and face-swap call as a subprocess rather than importing either runtime in-process.

Local Voice Synthesis Pipeline

Python (standard library only) · Local TTS

- Built a dependency-free, standard-library-only CLI that turns arbitrary-length text into one stitched audio file using a locally-hosted voice-cloning server, enforcing a hard loopback-only network restriction in code so text and audio never leave the machine.

Skills

AI & ML: LLMs, RAG, Prompt Engineering, Generative AI, AI Workflow Design, OCR Pipelines, Cognitive Services, Agentic Pipelines, Multi-agent Orchestration, AI Detection & Humanization

AI-Directed Development: Claude (Anthropic) - agentic orchestration, headless CLI integration as a production LLM backend, Google Gemini API - production integration, content & scoring pipelines, GPT-4/GPT-5 (OpenAI) - prompt engineering, structured extraction, GitHub Copilot - led org-wide rollout and adoption at Infosys, AI-Directed Full-Stack Delivery - ships production systems by directing frontier models as build partners, not unaided manual coding, Multi-Provider LLM Routing & Fallback Design - cost-ordered failover across providers so pipelines stay alive on free/low-cost tiers

Engineering: TypeScript, Next.js, Java, Python, RESTful APIs, Microservices, System Architecture, SOLID Principles, BeautifulSoup, Playwright, ReportLab, WebSocket

Cloud & DevOps: Azure DevOps, Docker, Kubernetes, CI/CD, Site Reliability Engineering, FastAPI, SQLite, Concurrent Threading

Data: MS SQL, T-SQL, Tableau, ETL, ERWIN Data Modeling, Row-Level Security, Multi-tenant Architecture, Fuzzy Matching, YAML-based Config Systems

Product: Roadmapping, OKRs, RICE, SAFe, Agile/Scrum, User Research, GTM Strategy, Retention & Churn Reduction, Stakeholder Management, Market Intelligence Systems

Leadership: Solution Architecture, Executive Communication, Responsible AI

Education

Florida International University, M.S. Engineering Management – Miami, FL, USA

GITAM University, B.Tech Mechanical Engineering – Hyderabad, India

Certifications

- Product Management · Infosys
- Software Product Management · University of Alberta, Canada
- Business Analytics · Wharton School, University of Pennsylvania