

SANDHYA H

Staff Software Engineer · AI Systems Architect · Generative AI Platform Builder

✉ sandhyah23@gmail.com · 🔗 [linkedin.com/in/sandhyahonnappa](https://www.linkedin.com/in/sandhyahonnappa) · 🌐 sandhyahonnappa.dev

📍 Bangalore, India · 📞 +91-9663075857

Professional Profile

Staff Software Engineer with 15+ years of experience, including 4+ years specializing in enterprise Generative AI platform architecture.

Proven track record as the key architect of JLL GPT — a multi-provider enterprise GenAI platform (chat, custom AI assistants, RAG, prompt management) serving ~45,000 monthly users globally — with end-to-end ownership across AI infrastructure, full-stack engineering, and cloud architecture.

Deep expertise in LLM orchestration, agentic AI (MCP), RAG pipelines, and AI gateway design, with Python backend experience spanning production systems (Flask API development, LiteLLM gateway) and personal AI projects (FastAPI, PostgreSQL, pgvector) — bringing a cross-stack perspective to AI system design.

Core Competencies

Domain	Skills
AI & GenAI	LLM Integration (Azure OpenAI, AWS Bedrock, Gemini, Baidu Ernie), RAG Pipelines, Agentic AI, MCP, LiteLLM AI Gateway, Prompt Caching, Prompt Engineering
Cloud & Infra	Azure (AKS, Cosmos DB, Front Door, APIM, App Service, Container Apps), AWS (S3, Lambda, EC2, RDS), Azure DevOps CI/CD
Backend	Python, FastAPI, Flask, ASP.NET Core 8, C#, Clean/Onion Architecture, CQRS (MediatR), SignalR, REST APIs, Node.js/Express
Data	PostgreSQL, pgvector, Azure Cosmos DB, Redis, Qdrant (Vector DB), Datadog, Splunk, Azure Monitor

Domain	Skills
Architecture	System Design, ADR Authoring, Multi-Provider Strategy Pattern, AI Gateway Design, Resilience Patterns, Security & Auth
Frontend	React 18, TypeScript, Next.js, Angular, Redux, Design Systems, Micro Frontend, PWA

Work Experience

Staff Software Engineer | JLL (Jones Lang LaSalle)

Apr 2022 – Present · Bangalore

JLL GPT — Enterprise Generative AI Platform *(Key Architect & Lead Engineer)*

- Owned end-to-end AI architecture of JLL GPT, a full enterprise GenAI web application used by 45,000 employees globally; designed the multi-provider routing strategy (Strategy Pattern + ChatBackendServiceLocator) enabling seamless LLM hot-swapping with zero code changes.
- Designed and built the Assistant module, letting users create custom AI assistants with configurable system instructions, model selection, and access management, grounded on multi-source knowledge (document uploads, Confluence, SharePoint links) indexed via an internal enterprise RAG service.
- Designed and built the Prompt Library and Collections feature for saving and sharing reusable prompts across business units.
- Architected the backend with Clean/Onion Architecture, CQRS, JWT/Okta auth, and Azure Cosmos DB; delivered real-time persistent chat via SignalR with full conversation management.
- Deployed and own the LiteLLM AI Gateway (Python) on AKS across two Azure regions with exact-match and semantic caching (Qdrant); designed Azure networking topology (Front Door → APIM → Container Apps → Private Endpoints) with rate-limiting and idempotency protections.
- Architected an MCP (Model Context Protocol) Gateway on Node/Express enabling agentic tool-use; authored ADRs and Claude Code setup guides; mentored engineers on AI architecture and resilience design patterns.
- 🏆 CoE Shining Star Award

- 🏆 JLL GPT - 2024 Stevie Award Winner — Bronze - Technical Innovation of the Year

LeaseLens AI (*Backend Contributor; Frontend Architecture Advisor*) — Built Flask API routes for early-stage document upload and LLM-interaction features; advised on frontend architecture (conversation persistence, document handling, prompt management).

Outlook Assistant AI (*Frontend Architecture Advisor & Contributor*) — Advised on UX and reviewed the initial release of this Office 365 add-in for email summarization, Q&A, and AI-drafted replies, reducing defects and improving adoption; contributed bug fixes and content-rephrasing features.

Specialist Front End Developer | L&T Technology Services (LTTS)

2019 – 2022 · Bangalore

- Led end-to-end architectural design across multiple React and Angular projects, including an IoT Web Application and an enterprise Design System (reusable component library).
- Acted as technical consultant for code reviews, performance troubleshooting, and BDD testing adoption; championed Storybook for component documentation.
- 🏆 Stellar Team of the Year 2019 · Outstanding Performer (Client Award)

Lead Front End Developer | SLK Software

2015 – 2019 · Bangalore

- Led a 15-person team building PayAdmin, a large-scale Insurance and Investment Plan System (Angular 7, TypeScript, Ag-Grid), migrating a legacy payroll system to a modern SPA with improved UX and performance.
- 🏆 Star Performer of the Year 2018–2019

Earlier Career — Bangalore (2010 – 2015)

- **Senior Full Stack Developer, Xerox** (2014–2015) — Built and optimized ASP.NET/C# web forms; resolved data integrity issues via SQL root-cause analysis.
- **Senior Software Developer, Tricon Infotech** (2012–2014) — Upgraded product to ASP.NET MVC; redesigned UI with responsive design principles.

- **Software Developer, Apex Decisions** (2010–2012) — Maintained and enhanced ASP.NET applications across the product lifecycle.

Key Projects

Nudge — AI Habit Tracker (*Personal Project · Python Backend*)

Full-stack app built with FastAPI, PostgreSQL, pgvector, LiteLLM, and MCP. Features a multi-repo architecture (nudge-mcp-server, nudge-agent, nudge-pwa), JWT auth, RAG-ready vector storage, and an MCP server for agentic habit logging. Hosted on Railway with managed Postgres. Serves as both a functioning personal tool and a portfolio demonstration of Python AI backend patterns.

Certifications & Learning

- AWS Certified Developer – Associate (DVA-C02) · In Progress (2026)
- Model Context Protocol (MCP) — JSON-RPC transport, tool/resource/prompt primitives, agentic patterns
- Claude Agent SDK — query() patterns, autonomous ADR writer and resilience auditor POCs
- LiteLLM AI Gateway — Production AKS deployment, semantic caching, cost governance
- Azure Architecture — AKS, Front Door, APIM, Private Endpoints, Cosmos DB

Education

Bachelor of Engineering — Computer Science

Visvesvaraya Technological University, Bangalore · 2006 – 2010

Achievements

- 🏆 CoE Shining Star Award — JLL
- 🏆 JLL GPT - 2024 Stevie Award Winner — Bronze - Technical Innovation of the Year
- 🏆 Outstanding Performer — Client Award, LTTS, 2019
- 🏆 Stellar Team of the Year — LTTS, 2019

- 🏆 Star Performer of the Year — SLK Software, 2018–2019

| *Open to AI Architect / Staff Engineer roles · Remote Roles*