

# Jayaprakash M

Senior Development Engineer (SDE II - AI & Systems)

+91 9791451155 | [mjpv1.ai@gmail.com](mailto:mjpv1.ai@gmail.com) | [jp-profile-webpage.netlify.app](http://jp-profile-webpage.netlify.app) | [linkedin.com/in/jayaprakash-murugan-7528b7229](https://linkedin.com/in/jayaprakash-murugan-7528b7229) | [github.com/mjpv1-ai](https://github.com/mjpv1-ai)

## PROFESSIONAL SUMMARY

**Senior Development Engineer (SDE II)** with a unique technical trajectory: combining **8 years of entrepreneurial hardware/software business experience** that built rigorous end-to-end system ownership, followed by **2.5 years of professional software engineering** focused on scalable backend orchestration and multi-agent voice systems. Specializing in high-performance distributed microservices, low-latency conversational architectures, and emerging GenAI/LLM-driven engineering. Highly enthusiastic about bringing hands-on expertise in automated tool routing, isolated stream services, and custom protocol design to a high-impact infrastructure team.

## TECHNICAL SKILLS

**AI Architecture & Systems:** GenAI/LLM-Driven Solutions, Multi-Agent Systems, Agent-to-Agent (A2A), Model Context Protocol (MCP), LiveKit STT/TTS Pipelines, Custom A2A Orchestration, Distributed Microservices.

**LLM Training & Inference:** Reinforcement Learning (RLHF, GRPO, PPO), Quantized Model Deployment, RAG, Inference Optimization.

**Programming Languages:** Rust, Python (Async/OOP), Go (Golang), C/C++, Lua, SQL.

**Infrastructure & Cloud:** Cloud Technologies (AWS, GCP), Docker, Custom Networking (OpenWRT, mTLS), P2P/NAT Traversal.

## APPLIED AI RESEARCH & ARCHITECTURAL INNOVATIONS

### LiveKit A2A Relay | *Rust Open Source Crate*

- Designed and built an ultra-low latency bridge (`livekit-a2a-relay`) in Rust connecting LiveKit WebRTC streams to A2A-compliant endpoints.
- Implemented custom jitter buffering, RMS-energy Voice Activity Detection (VAD), and turn management to handle user interruptions and guarantee stutter-free playback.

### LLM Reasoning Optimization (Gemma 3) | *Reinforcement Learning & Finetuning*

- Applied Group Relative Policy Optimization (GRPO) to train and finetune the Gemma 3 LLM, significantly enhancing its multi-step reasoning capabilities and optimizing context management.
- Curated specialized datasets and developed robust output safety parsers, earning two Bronze Medals in global Kaggle machine learning competitions.

### Autonomous Multidrone Trajectory Planning via DRL & Blockchain | *M.Tech Thesis Research*

- Formulated a DRL (PPO) architecture for multi-drone navigation in unknown, partially observable environments, seamlessly transitioning from spatial exploration to tactical exploitation.
- Engineered a 3D array and segment tree data structure to dynamically optimize spatial range querying, drastically reducing the required exploration iterations from 150,000+ episodes to under 500.
- Designed a decentralized, lightweight blockchain network to secure inter-drone communication, utilizing SHA256 hashing to enable trustworthy, multi-agent sharing of explored strategic paths.

### MedSwarm - Decentralized Agentic Intelligence | *Multi-Agent Swarm Architecture*

- Developed a decentralized, multi-agent system utilizing localized, quantized MedGemma models to provide highly auditable clinical reasoning and planning in low-connectivity environments.
- Implemented persona-based adversarial agent debating to act as a rigorous learning and evaluation loop, stress-testing complex decisions and minimizing hallucinations prior to execution.

### ProdAlve - Generative AI Observability | *Personal Hackathon Project (Top 10 Finalist)*

- Developed a serverless Generative AI platform for real-time anomaly detection, utilizing **Vertex AI** and **Gemini 2.5 Flash** to optimize autonomous agent decision-making.
- Built an interactive telemetry frontend to visualize real-time data streams and make the internal reasoning and multi-step planning of the AI agents fully auditable to human operators.

## PROFESSIONAL EXPERIENCE

### Truminds Software Systems

*Principal AI & Systems Engineer (SDE II)*

Bengaluru, India

Aug 2023 – Present

- Designed and implemented isolated LiveKit services within a real-time voice streaming pipeline, handling continuous audio ingestion and processing across independent STT and TTS layers.
- Reengineered the traditional core architecture by replacing the standard foundational LLM module with a custom-designed, low-latency A2A orchestrator built to serve entirely as the central reasoning, execution, and tool-routing driver.
- Refactored legacy FastAPI orchestration engines in **Python** into highly scalable, disaggregated microservices, reducing average system response latency by **60%** for on-premise and edge inference deployments.
- Established rigorous evaluation frameworks and dynamic DB schema parsing, transforming black-box model outputs into structured, trustworthy logic that human operators can audit and act upon.

### Truminds Software Systems

*Backend & Firmware Developer (Distributed Systems & Networking)*

Bengaluru, India

Aug 2023 – Present

- Developed a robust session management backend in **Go (Golang)** capable of handling concurrent P2P streams with minimal latency, resolving critical NAT traversal connectivity challenges.
- Coded optimized **C++ firmware** for embedded devices to execute real-time MQTT commands, leading the migration to ECC-based certificates (mTLS) for deployment security.

## EDUCATION & ACCOLADES

---

**Indian Institute of Information Technology, Design and Manufacturing**

*Master of Technology in Computer Science and Engineering (Data Science & AI)*

Kancheepuram, India

2023

**Professional Qualifications:** GATE-qualified professional (CS/IT & Geomatics Engineering).

**Professional Certification:** A2A: The Agent2Agent Protocol (DeepLearning.AI, 2026).

**B.E. in Computer Science:** Tamilnadu College Of Engineering, Coimbatore (2013).