# ADARSH KUMAR SINGH

+91-94394-40544 | adarshentity098@gmail.com | linkedin.com/in/i-am-entity | cyber-duelist.github.io |

Ghaziabad, Uttar Pradesh

## SUMMARY

Applied AI Engineer and Solutions Architect specializing in enterprise-grade agentic workflows, real-time AI infrastructure, and deploying foundation models into production. Proven track record of building complex multi-agent systems, highly accurate RAG pipelines, and low-latency voice streaming architectures using Python, FastAPI, and Django. Adept at translating complex client requirements into scalable deployments on AWS and Kubernetes, while implementing robust LLM guardrails for strict enterprise compliance. Oracle GenAI and Data Science certified.

## TECHNICAL SKILLS

Multi-Agent Orchestration | ReAct Framework | Retrieval-Augmented Generation (RAG) | Voice AI (TTS, STT) | Large Language Models (LLaMA 3, GPT) | Enterprise Deployments | Python | FastAPI | Django / DRF | WebSockets

## WORK EXPERIENCE

### Software Incubator (AKGEC)

Jan 2026 - Present

**Technical Lead & AI Engineer**

- Partnered with internal stakeholders to translate complex compliance requirements into scalable AI evaluation frameworks and concrete technical specifications.
- Architected and deployed a stateful enterprise compliance agent with persistent memory, tool-calling capabilities, and dynamic routing (70B to 8B models), ensuring 99.9% uptime.
- Built an LLM-as-a-Judge Evaluation Dashboard scoring agent responses across 5 dimensions, and implemented pre-execution Input Guardrails blocking 100% of adversarial injections.

### Programming Club (AKGEC)

Aug 2024 - Dec 2025

**Engineering Lead (RAG & AI Systems)**

- Architected a highly accurate RAG pipeline utilizing ChromaDB for large-scale PDF ingestion, reducing LLM hallucination by 78% via source-grounded generation and citation tracking.
- Designed and exposed high-throughput FastAPI REST endpoints for production usage, implementing API rate limiting and structured JSON responses.
- Led cross-functional collaboration and mentored a team of 15+ engineers on Python best practices, Git workflows, and robust API development.

### Cloud Computing Club (AKGEC)

Aug 2023 - Jul 2024

**Cloud AI Deployment Engineer**

- Managed end-to-end deployment of 5+ containerized AI microservices utilizing Docker and Kubernetes across AWS and OCI environments.
- Streamlined infrastructure operations by automating 3 CI/CD pipelines via GitHub Actions, accelerating deployment velocity by 60%.

## PROJECTS

### Real-Time Voice AI Agent WebSockets

Developed a low-latency, bidirectional voice agent utilizing WebSockets and Server-Sent Events (SSE) for real-time speech-to-text (STT) and text-to-speech (TTS) streaming. Optimized pipeline latency to achieve sub-500ms time-to-first-byte (TTFB), simulating enterprise-grade voice interactions for customer service scenarios.

Python, FastAPI, WebSockets, TTS/STT, LLMs

### Autonomous AI DevOps Swarm ReAct

Engineered a 4-agent swarm using ReAct loops to autonomously detect, debug, and remediate CI/CD pipeline failures via sandboxed execution. Reduced mean-time-to-recovery (MTTR) by 90% and achieved an 85% first-pass resolution rate.

Python, LLaMA 3, Pytest, ReAct Framework

### AI Code Review & Risk Assessment Django

Built a robust Django REST application processing 100+ GitHub PRs, leveraging LLMs to detect edge-case logic bugs and vulnerabilities missed by static analysis. Integrated a supervised ML pipeline (Random Forest) to classify document risk.

Django, REST Framework, Groq, Machine Learning

## EDUCATION

**B. Tech in Electronics and Communication Engineering (CGPA: 7.9)** Aug 2022 - May 2026

Ajay Kumar Garg Engineering College (AKGEC)

## CERTIFICATIONS

| | |
|---|---|
| Oracle GenAI Certified Professional | 2025 |
| Oracle Data Science Professional Certification | 2025 |
| IoT and Industrial Automation Certification | 2025 |

## ADDITIONAL SKILLS

**Frameworks:** FastAPI, Django / Django REST Framework, WebSockets, Pytest, ChromaDB

**Cloud / DevOps:** Docker, Kubernetes, AWS (EC2, S3), Oracle Cloud (OCI), CI/CD, GitHub Actions

**Tools:** Git, REST APIs, Server-Sent Events (SSE), LLM-as-a-Judge