

# NADEEM SAJJAD

Senior AI Engineer · Forward-Deployed Engineer

Agentic AI Systems, LLM Orchestration & End-to-End Client Delivery · 12+ Years

Remote (Americas overlap) · +91 88265 17429 · nadeem.sajjad.1991@gmail.com

[github.com/nadeemcrite](https://github.com/nadeemcrite) · [linkedin.com/in/nadeem-sajjad](https://linkedin.com/in/nadeem-sajjad) · [devnadeem.com](https://devnadeem.com)

## SUMMARY

AI engineer with 12+ years shipping production systems, now focused on agentic AI delivered end to end for clients. As a forward-deployed engineer through Invisible, I embed as the sole dedicated engineer, translate client requirements and SOWs into architecture, and own delivery through to production. My work spans multi-agent systems, LLM orchestration, and RAG, paired with the production discipline these systems need: automated tests, error handling, retries, and monitoring. Comfortable across external integrations (APIs, webhooks) and real-time and voice tooling (ElevenLabs, WebRTC). Open-source contributor to LangChain and LangFlow.

## PROFESSIONAL EXPERIENCE

### Invisible Technologies

Sept 2021 – Present

Senior Software Engineer / Team Lead (Forward-Deployed Engineer)

Remote

Embedded as the sole dedicated engineer across multiple enterprise clients. Translated requirements and SOWs into architecture and owned delivery end to end, from prototype to production launch. Led a small team of two engineers.

- **Built the Ontology Platform**, an end-to-end agentic system that turns a client's raw database into queryable, ontology-driven data products. Cut data-modeling effort from weeks to under a day across 100+ unstructured tables.
  - *Multi-agent generation with HITL*: a multi-agent system builds the ontology step by step on a fully automated path; users step in through chat to correct a join or generated SQL whenever they are not satisfied.
  - *AI Data Dictionary*: enriches every table and column and indexes them in Milvus, making the whole schema semantically searchable in plain language.
  - *Ontology Studio*: users design and refine their own ontology (add, remove, and combine entities), then map it to the dictionary to generate real, runnable pipelines on their own data.
  - *Orchestration and serving*: Temporal runs the heavy jobs, and finished ontologies are served as an MCP so clients and stakeholders can query their own data directly.
- **Designed a unified document-ingestion pipeline** that parses heterogeneous client inputs (PDF, PPT, Excel) with local LLMs and OCR, then structures and maps the data into 10+ client-specified report formats. Cut report turnaround from ~2 days to under 30 minutes.
- **Shipped to production quality**: automated tests (unit and end-to-end), error handling and retry strategies, and monitoring across a Kubernetes-based delivery stack (Helm, Argo, Docker, GitHub Actions), cutting client infrastructure costs by 25%. Hardened reliability with fault-tolerant pipelines and load testing at 10K concurrent users.

### 42 Students

Dec 2019 – Sept 2021

Lead Engineer and Co-Founder

Bengaluru, India

- Architected and shipped a Learning Management Platform (WYSIWYG editor, plagiarism detection, multi-platform support, CI/CD, Kubernetes autoscaling) serving 25K+ students across institutions.
- Built a college Ranking Platform surfacing intake, placement, and NIRF metrics, generating hundreds of qualified admissions leads per month and a new education-loan revenue stream via NBFC partners.

### Acko General Insurance

Dec 2017 – Nov 2019

Software Engineer

Bengaluru, India

- Built the Acko Data Platform, integrating BigQuery and Tableau for real-time dashboards; led a team of 5 and delivered in one month.

### Easeit Technologies

Feb 2014 – Nov 2017

Software Engineer & Co-Founder

New Delhi, India

- Co-founded the studio and led a team of 3, delivering 30+ production web apps (PHP, MySQL, WordPress, Angular).

## TECHNICAL SKILLS

**AI & LLM:** Agentic AI, Multi-Agent Systems, Tool Use / Function Calling, LLM Orchestration, RAG, Prompt Engineering, MCP, Multi-Turn Conversational & HITL Workflows, Vector Search (Milvus), Voice & Real-Time (ElevenLabs, WebRTC)

**Engineering & Reliability:** Automated Testing (unit, end-to-end), Error Handling & Retry Strategies, Load Testing, Monitoring, API & Webhook Integrations

**Languages & Frameworks:** Python, TypeScript, JavaScript, Next.js, React, Node.js, Django, FastAPI

**Data & Infrastructure:** PostgreSQL, Redis, BigQuery, Milvus, Temporal, Kubernetes, Helm, Argo, Docker, GitHub Actions