

AYUSH RANJAN ROY

AI Engineer | Multi-Agent Systems, RAG, LLM Infrastructure & Cloud Platform

Gurugram | royayush0611@gmail.com | +91-9540968483 | linkedin.com/in/ayush-ranjan-roy | github.com/AyushRanjanRoy-01

Claude Certified Architect - Foundations (Anthropic, May 2026) | ayushranjanroy.com

SUMMARY

AI Engineer contributing to production multi-agent and RAG systems (LangGraph, Claude, Azure AI, MCP, HITL), with 2 years modernizing Accenture MxDR / Adaptive MxDR (AMxDR) AWS cloud architecture. On large multi-engineer programs: implemented inference, routing, tool gateways, and observability components with measured accuracy/latency gains - not sole ownership of the full enterprise O2C/D&D product.

EXPERIENCE

Full Stack LLM Development Analyst / AI Engineer

Aug 2025 - Present

Accenture - GenAI & Data, Gurugram | Multi-Agent Systems (team contributor)

- Contribute multi-agent components for enterprise Order-to-Cash / Dispute & Deduction programs using LangGraph, durable state, and human-in-the-loop gates (one engineer among a larger delivery team).
- Implemented async multi-model inference (Claude Sonnet, GPT-4o/5, Azure Document Intelligence) with ensemble/confidence routing and hot-reload model gateway config.
- Measured in controlled runs: classification accuracy ~+40%; batch capacity 9->50 pages; multi-page latency ~-55% / single-page ~-30%; paths supporting 300+ docs/day at 97%+ accuracy on the program.
- Built MCP-style governed tool access and operator-facing agent observability (OpenTelemetry + Datadog); hardened multi-agent failure paths.
- Authored architecture notes and ran LangGraph/LangChain KT for 4 engineers (~60% faster ramp-up).

Platform Engineer

Aug 2023 - Aug 2025

Accenture - Security | MxDR / Adaptive MxDR (AMxDR) AWS Cloud Platform

- Led platform engineering work modernizing MxDR -> AMxDR AWS cloud architecture: Terraform, GitHub Actions CI/CD, production incidents across ~20 microservices / 800+ enterprise accounts (~35% faster deploys, ~30% infra cost reduction in measured work).
- Owned Redis/ElastiCache shared metadata services; Terraform across 6 environments; modernized 8+ Lambdas (Python 3.8->3.12) cutting runtime failures ~50%.
- Automated Secrets Manager + IAM migrations; hardened 12+ Java/Python services; restricted-egress proxy automation; CodeQL/Dependabot remediation.
- mTLS/OpenTelemetry 5-min synthetic health into enterprise SIEM; blue-green (Route 53/ASG) with zero customer-facing downtime on a live pipeline upgrade.

OPEN SOURCE

HelixOps (github.com/AyushRanjanRoy-01/HelixOps) | RAGGym (rag-gym.streamlit.app)

- HelixOps: multi-agent RCA + RAG over runbooks, human-gated remediation, JWT/RBAC, decision traces.
- RAGGym: swappable RAG stack, self-correcting retrieval, RAGAS eval, AI-graded practice mode.

SKILLS

Agents/GenAI: LangGraph, LangChain, Multi-Agent, MCP, HITL, RAG, Claude, GPT, Azure AI Foundry, ensemble routing, RAGAS
Platform/Security: AWS (Lambda, EC2, S3, ASG, Route 53), Terraform, Secrets Manager, IAM, mTLS, SIEM, Docker, GHA, CodeQL
Data/Obs: PostgreSQL, Redis, pgvector, Qdrant, FastAPI, async Python, OpenTelemetry, Datadog, CloudWatch

EDUCATION

B.Tech ECE - VIT Chennai

2019-2023

Coursework: AI; ML essentials; Cryptography & Network Security; InfoSec Audit; DSA; Operating Systems.