

# Mohideen Hajamohideen

Senior Solutions Architect · GenAI & Applied AI · Multi-Cloud · Agentic Systems

Mohideen.HajaMohideen@outlook.com | [linkedin.com/in/mohideenabducader](https://www.linkedin.com/in/mohideenabducader) | [github.com/mohideen-aws](https://github.com/mohideen-aws) | London, UK

## PROFILE

Senior Solutions Architect with **14+ years** designing and shipping production systems for enterprise - the last four years focused on **GenAI and LLM-powered applications** on AWS Bedrock, SageMaker, Azure, and GCP. At AWS Professional Services I operate as the post-sale technical lead for complex AI engagements across EMEA: embedded with customer engineering and business teams, writing production code, guiding architecture, and translating between business outcomes and engineering trade-offs. Proven across four archetypes - **Applied AI / RAG engineering, cloud & platform architecture, forward-deployed delivery, and adoption & change management**. Equally comfortable benchmarking foundation models, debugging agent routing across microservices, or sitting with end users for two days before writing a single line of code. I build trusted-advisor relationships with technical and business decision makers, systematise delivery patterns into reusable accelerators, and mentor a 12-person EMEA practice. **Multi-cloud production experience** (AWS primary · SA-Pro; ~4 yrs Azure; hands-on GCP); Python daily driver.

## CORE COMPETENCIES



## EXPERIENCE

### Amazon Web Services - Professional Services, EMEA

Jul 2021 - Present

#### Senior Delivery Consultant - GenAI & Cloud

London, UK · Promoted to Senior Q4 2024

- ▶ Post-sale technical lead for enterprise GenAI & cloud engagements across financial services, legal, retail, and tech - **embedded with customer engineering and business teams**, setting architecture, writing production code, and guiding customers from pilot to operational adoption.
- ▶ **Contract Intelligence on AWS Bedrock (production, in daily use)**: designed document-aware chunking for long-form legal contracts, hybrid retrieval with re-ranking, and **benchmarked Claude Opus, Amazon Nova, and Google Gemini** for retrieval and answer quality. **~73% reduction in contract review time (~2.5 hrs → ~40 min)** - driven by two days spent with the legal reviewers mapping their real workflow *before* writing a single line of code.
- ▶ **Sub-millisecond latency PoC on AWS Outposts** for a trading firm: owned bare-metal networking, kernel tuning, and the end-to-end benchmark harness. Customer signed a multi-year deal on the back of the result. Awarded *Bias for Action Q4 2023* for delivering ahead of deadline.
- ▶ **Multi-agent agentic workflows**: ReAct-style reasoning, tool selection and routing, self-reflection loops, risk-tiered HITL approval gates, connected to legacy enterprise systems (SharePoint, VPC-isolated infrastructure, SAML/OAuth). Debugged agent logic failures by tracing conversation IDs across distributed microservices under live load.
- ▶ **Delivered Citrix-to-Amazon AppStream migration for ~1,000 users** - cut desktop delivery cost by **35%**; freed budget was later redirected into customer AI initiatives.
- ▶ **Built internal QuickSight FinOps dashboard** adopted across the EMEA region; hosted an internal webinar with **122 attendees** on the patterns behind it.
- ▶ **Spoke at AWS Summit London** (300+ audience) on *Geospatial Foundation Models for Sustainability Monitoring*. Won "**Best Show**" at an internal sustainability hackathon for a GenAI tool cross-referencing satellite imagery with ESG filings to flag potential greenwashing.
- ▶ **Systematised delivery patterns** into reusable accelerators - RAG reference architectures, evaluation frameworks, ACM workshop guides, agentic workflow blueprints - adopted across a 12-person EMEA practice (5 × L5, 7 × L6) and fed back to product leadership.
- ▶ **Hiring and mentoring**: conducted 20+ interviews for architect and consultant roles annually; mentored L5/L6 consultants on GenAI and cloud engagements.

### HCL Technologies

Aug 2017 - Jul 2021

#### Associate Consultant - Multi-Cloud Platform Engineering

London, UK · embedded with global FMCG client

- ▶ **Four years of Azure production** as embedded multi-cloud consultant for a global FMCG client running workloads across **AWS, Microsoft Azure, and GCP simultaneously**. Direct hands-on with Azure (VMs, storage, networking, Azure DevOps), GCP, and AWS in the same week.
- ▶ Shipped **30+ AWS Service Catalog products** for self-service infrastructure provisioning - most backed by Python Lambda functions - eliminating ticket-queue waits for developer teams.
- ▶ Hands-on multi-cloud IaC with **Terraform and Ansible** across AWS, Azure, and GCP production environments. Built Prometheus / Grafana / CloudWatch monitoring stack for 24/7 operations.
- ▶ Collaborated directly with AWS, Microsoft, and Google Professional Services on joint customer engagements. Delivered executive-level stakeholder workshops and cross-cloud governance reviews.
- ▶ **Excellence Award × 3** - three consecutive quarters from a banking client (a first for that engagement).

### Wipro Technologies

Nov 2010 - Aug 2017

#### Senior Technical Consultant - Infrastructure Automation

- ▶ Authored a **PowerShell patch-compliance automation** that replaced ~120 hours of manual work per month - finance-calculated savings **>\$1M/year**.
- ▶ Owned DR drills, Exchange migrations, and the Windows/VMware infrastructure fundamentals that still shape how I think about reliability and incident response on production AI systems today.

## SELECTED PROJECTS

### Self-Evolving Agentic Systems - Human-like Continual Learning

R&D · Python · LangGraph

Research track exploring agents that **reflect, remember, and rewrite themselves**: episodic + semantic memory layers, reflection loops that critique prior runs, and a meta-controller that evolves the agent's own tools and prompts against measured outcomes. Targets the core limitation of current agent stacks - they don't learn from yesterday.

[Python](#) · [LangGraph](#) · [Claude Opus](#) · [Reflection Loops](#) · [Episodic Memory](#) · [Vector Stores](#) · [Tool Self-Authoring](#)

### Secure-OpenClaw - Agentic AI Security Framework [Open Source](#) · [Python](#)

Production-grade agentic AI framework: **five security layers**, local LLM integration (Ollama), **risk-tiered HITL approval gates** with timeout safe-state, Docker sandboxing, and multi-channel alerting. A public, working demonstration of the safety and adoption patterns that make enterprise AI deployments stick.

[Python](#) · [Docker](#) · [Ollama](#) · [Agent Orchestration](#) · [HITL](#) · [Defense-in-Depth](#) · [ReAct Patterns](#)

### MetaCalib - LLM Metacognition Benchmark [Kaggle / DeepMind](#) · [\\$200K Competition](#)

Novel LLM calibration benchmark for the Kaggle / Google DeepMind *Measuring AGI Progress* competition. Tests model confidence calibration (Expected Calibration Error), knowledge-boundary detection, and error self-detection - the evaluation rigour enterprise AI adoption programmes need to build stakeholder trust.

[Python](#) · [kaggle-benchmarks SDK](#) · [ECE](#) · [LLM Evaluation](#) · [Benchmark Design](#)

### Greenwashing Detection - Satellite x ESG Cross-Reference [Hackathon](#) · ["Best Show"](#)

GenAI tool that cross-references satellite imagery against public ESG filings to flag potential greenwashing. Built in 48 hours; won *Best Show* at the internal AWS sustainability hackathon. Geospatial foundation models + LLM reasoning in production-adjacent form.

[Bedrock](#) · [Geospatial Foundation Models \(Clay / Prithvi\)](#) · [Sentinel-2](#) · [LangChain](#) · [Python](#)

## EDUCATION

---

**BSc Chemistry** (with Mathematics & Physics) - *First Class Honours*

[Madurai Kamaraj University](#)

## CERTIFICATIONS

---

AWS Certified AI Practitioner

[Amazon Web Services](#)

AWS Certified Solutions Architect - Professional

[Amazon Web Services](#)

AWS Certified Solutions Architect - Associate

[Amazon Web Services](#)

FinOps Certified Practitioner

[FinOps Foundation](#)

## TECHNICAL SKILLS

---

**Languages:** Python (daily driver - Boto3, FastAPI) · JavaScript / Node.js · Bash · PowerShell

**GenAI / LLMs:** RAG pipelines · Chunking strategies · Embeddings & vector search · Prompt engineering · Agentic workflows · Function calling · Evaluation loops · Guardrails · Model benchmarking

**LLM Platforms:** AWS Bedrock (Claude Opus, Amazon Nova) · Amazon SageMaker · OpenAI (GPT-4, Assistants) · Google Gemini · Ollama · LangChain · LangGraph

**Retrieval & Vector:** Amazon OpenSearch · Pinecone · Hybrid search · Re-ranking · Knowledge-base design

**Cloud:** AWS (primary, SA-Pro) · Microsoft Azure (~4 yrs production) · Google Cloud Platform · Multi-cloud production delivery

**Infrastructure:** Terraform · Ansible · Kubernetes · Docker · Lambda · CI/CD · AWS Service Catalog · AppStream · Outposts

**Observability & FinOps:** Prometheus · Grafana · CloudWatch · QuickSight · Cost optimisation & CUR analytics

**Ways of working:** Customer workshops · Exec-level presentations · Architecture reviews · Roadmap planning · Post-sale advisory · Hiring & mentoring · Adoption & Change Management