



Sydney, NSW, Australia

(+61) 415 101 418

daniyalxmalik@gmail.com

daniyalmalik.com

linkedin.com/in/malikdaniyalahmed

github.com/daniyalmalikoo7

Daniyal Malik

Senior Software Engineer & GenAI Architect
AWS Certified Generative AI Developer — Professional

Senior engineer and architect with **7+ years shipping production systems end-to-end** — across startups, enterprise, and consultancy. A **MERN / TypeScript full-stack engineer who went deep on both ends**: distributed backends (Node microservices, Kafka, BullMQ, SQL/NoSQL, event-driven, n-tier, serverless) and modern frontends (**Next.js App Router, React, RSC**) with real UI/UX and **SEO / GEO** sense. I've led enterprise platforms to **40,000+ users at 99.5% uptime** and integrated modern stacks into **legacy enterprise systems** (.NET Core, Sitecore). As an **AWS Certified Generative AI Developer (Professional)**, I architect to the Well-Architected GenAI Lens and ship production **multi-agent RAG** and voice AI (2.4M+ calls) — grounded, evaluated, and governed.

40,000+ USERS SERVED	99.5% PRODUCTION UPTIME	2.4M+ AI CALLS RUN	80% LATENCY CUT · 800→150MS	80+ AGENTS SHIPPED	7+ yrs IN PRODUCTION
--------------------------------	-----------------------------------	------------------------------	--	------------------------------	--------------------------------

EXPERIENCE

Preamble — Founder & AI Architect

2025 – Present · Sydney

Production multi-agent RAG platform for Australian construction tendering

Architected an **8-stage multi-agent pipeline** that reads 50–150-page tenders and produces **cited, priced quote drafts** — deterministic where it matters, LLM where it adds value, with every line verified to the character against its source page so **no hallucination ships silently**.

Owned the full **multi-tenant SaaS**: Claude orchestration, **pgvector** retrieval with reranking, LlamaParse + Voyage AI, **Langfuse** eval & observability, ground-truth eval sets, Stripe billing, Clerk auth — architecture, data model, agents, and UX.

007 AI — Co-Architect & AI Engineer

2024 – 2025 · Voice AI SaaS

Multi-tenant voice-AI platform — autonomous outbound calling at scale

Co-architected a platform where autonomous agents ran calls end to end — qualification, objection handling, booking, and CRM sync. **2.4M+ calls processed, with paying customers in production**.

Engineered **13 mid-call server tools**, 7 concurrent BullMQ queues, **Kubernetes** orchestration, Prometheus observability, timezone-aware DNC compliance, ElevenLabs + Twilio telephony, and Stripe metered billing on **AWS**.

Everest Engineering — Senior Software Engineer

Oct 2022 – 2025 · Melbourne (Remote-first)

IT consultancy — Technical Lead & Architect on enterprise client delivery

BUPA Wellbeing Dashboard — technical lead & architect, leading 3 senior engineers. Designed an **n-tier architecture for 2,000 concurrent users at <200ms**, serving **40,000+ users at 99.5% uptime** over an 8-month delivery. Cut recommendation-engine latency **80% (800ms → 150ms)** via PostgreSQL indexing and connection pooling.

Integrated a React / TypeScript frontend into **legacy enterprise systems via a BFF over .NET Core APIs and Sitecore CaaS**; event-driven Node microservices on **Azure** (App Gateway, Redis, Application Insights), GitHub Actions CI/CD, and **90% test coverage**. Impact: **-25% support tickets, +40% engagement**.

Education Horizons (EdTech) — API-first Node microservices plus a React Native app for **10,000+ students and parents** with real-time sync.

WeAreNova · Signup Solution · Aristostar — Full Stack Engineer

2020 – 2022 · UK / UAE (Remote)

Healthcare, B2B, and security SaaS across startups and product teams

Shipped TypeScript / Node / React SaaS: a **practitioner-matching platform for 500+ healthcare professionals** (WebSocket real-time), a **HIPAA-compliant medical-billing** system, and security / queue platforms — JWT & OAuth, RBAC, Redis sessions, Docker + Nginx.

Led an **Angular → React legacy modernization** of an authentication system, raising type safety and test coverage while keeping the product live.

Atrix Digital Solutions — Full Stack Developer (JS / TS)

2017 – 2020 · Part-time

Where my engineering career began

Built full-stack web apps end to end — **Angular / React** frontends and **Node.js / Express** backends over **SQL / NoSQL** databases — laying the foundation I've built on across every role since.

GENAI ARCHITECTURE ON AWS — CERTIFIED DEPTH

The 5 GenAI domains (certified)

FM integration & data — Amazon Bedrock, fine-tune vs RAG, enterprise data prep

Prompt & application design — prompt engineering, model chaining, orchestration

Agentic workflows & advanced RAG — agents, tools/APIs, vector retrieval

Model fine-tuning & optimization — domain adaptation tradeoffs

Security, governance & responsible AI — fairness, privacy, safe deployment

Well-Architected GenAI Lens — 6 pillars

Operational excellence — deploy, monitor KPIs, automate pipelines

Security — IAM, data protection, Bedrock Guardrails

Reliability — service quotas, hallucination & accuracy handling

Performance — FM selection, provisioned vs on-demand, latency

Cost optimization — token-based FinOps, measurable ROI

Sustainability — right-sized models, serverless footprint

RAG & DATA ENGINEERING

Retrieval & generation

Chunking strategy — fixed & semantic; **ontology / metadata-aware** retrieval

Pipelines & storage

Event-driven ingestion with **Kafka / BullMQ** and **Step Functions**

Hybrid retrieval with reranking; grounded generation with source citation

Quality via **ground-truth eval sets** & GenAI metrics — **faithfulness, context precision/recall, answer relevancy**

Human-in-the-loop review, governance, and deliberate **RAG-vs-fine-tuning** tradeoffs

Data structuring & verification via **AWS Glue / LlamaParse**

Right storage for the job — **blob (S3) vs relational vs vector vs cache**

SQL / NoSQL modeling, indexing, and connection pooling for scale

SELECTED PROJECTS

Claude Code Workflow Collection

github.com/daniyalmalikoo7

Open-source AI-native SDLC: **80+ production-grade agents across 4 repos** (SaaS build, DevOps/IaC on Terraform + AWS, codebase rescue, UI uplift) — spec-driven, phase-gated, with eval pipelines, hallucination guards, and Playwright visual QA.

Dan's Lab daniyalmalik.com

GenAI consultancy plus an autonomous design studio (generative media via Higgsfield) with a live "AI twin."

ProposalPilot [RFP](#) → [proposal](#) [RAG](#)

Next.js 16 · tRPC · pgvector RAG — converts RFPs into tailored, sourced proposals.

AI Code-Review Assistant [LLM](#) + [git hooks](#)

TDD-built TypeScript tool that brings LLM review into the git workflow.

TECHNICAL SKILLS

ARCHITECTURE

Microservices · event-driven · n-tier · BFF · monorepo · serverless · distributed systems · legacy integration · scale & performance

AWS

Bedrock + Guardrails · Lambda · ECS / Fargate · S3 · Step Functions · Glue · SageMaker JumpStart · Amazon Q · IAM · IaC (Terraform)

GENAI

Claude API · advanced RAG · multi-agent orchestration · LangChain / LangGraph · evals & guardrails · voice AI (ElevenLabs, Twilio) · Langfuse

BACKEND / DATA

TypeScript · Node.js · Python · FastAPI · tRPC · Kafka · BullMQ · PostgreSQL · MongoDB · MySQL · Redis · pgvector · Prisma

FRONTEND / UX

Next.js (App Router, RSC) · React · React Native · TypeScript · UI/UX · Core Web Vitals · SEO + GEO

CLOUD / DEVOPS

Azure (App Service, App Gateway, App Insights) · Kubernetes · Docker · CI/CD · Prometheus · OpenTelemetry

PRACTICE

Spec-driven development · TDD (90% coverage) · Agile / Scrum · code quality (SonarCloud)

EDUCATION & CERTIFICATION

AWS Certified Generative AI Developer — Professional

[verify on Credly](#) ↗

BSc, Electrical & Electronics Engineering

COMSATS Institute of Information & Technology · 2013–2017

Capstone: autonomous surveillance robot — image processing, obstacle avoidance, and face recognition with SMS alerting.