

Kyle Doidge

Senior Full Stack Engineer · AI Products

Leeds, UK · Remote (UK / EU) · 07913 348271 · kyle.blue.doidge@gmail.com
linkedin.com/in/kyle-doidge · github.com/kyle-blue · kblue.io · stillapp.io

PROFILE

Full stack engineer with 5 years of experience shipping products end to end, from RAG and agent pipelines through the backends that serve them to the interfaces people actually use. Founder of Still, an AI notes and assistant product built alone across Go, React Native and a fine-tuned open-weight model, and currently leading frontend delivery on an AI assurance platform for the Royal Navy under SC clearance. As comfortable owning a design system as a database migration or an eval harness. Looking for an AI product team small enough that individual ownership still counts.

EXPERIENCE

Founder & Engineer, Still (Still Group Ltd)

March 2026 – Present

stillapp.io · Company no. 17188025 · Remote

Sole founder, director and engineer of an AI-native notes and personal assistant product. Own every layer, from retrieval and the model pipeline through the backend and mobile client to infrastructure and product direction.

- Built the retrieval-augmented generation (RAG) layer the product runs on, covering chunking, embedding and semantic search over a user's notes, so the assistant answers from the user's own corpus rather than model priors and writes back into it
- Designed an LLM agent pipeline on Anthropic's API with tool use, multi-step orchestration and structured output, served from a Go backend to a React Native client
- Fine-tuning a Qwen-family open-weight model to run complex note-processing tasks at a fraction of frontier-model cost and latency, evaluated against a task set built from real usage traces
- Made a non-deterministic agent pipeline reliably regression-testable in CI, using a deterministic fake provider alongside testcontainers-backed integration tests with per-test user isolation
- Full Detox end-to-end coverage on React Native, chosen over Maestro after evaluating both on debugging ergonomics and CI stability

Full Stack Software Engineering Consultant, Methods (Alten)

August 2025 – Present

Remote · SC Cleared

Delivering software for UK defence and central government clients, across greenfield AI product development and modernisation of large microservice estates.

Royal Navy / MoD, AI Ethical Assurance Platform

- Joined a stalled project with no agreed product direction and a fixed delivery date; produced the Figma design system and end-to-end flows that became the team's single reference point, unblocking parallel work across four workstreams
- Led frontend delivery as the only dedicated frontend engineer, designing and building a bespoke SvelteKit design system of 40+ accessible components and then shipping the full application on top of it
- Built the platform's core capability, including a user-configurable assurance workflow and dynamic form engine, in-app chat, notifications, user management and audit logging
- Implemented role-based access control across 12 distinct service roles, spanning SvelteKit route guards, backend endpoint authorisation, permission-aware response shaping and supporting schema migrations
- Used AI-assisted and agentic development workflows where permitted to compress delivery timelines, holding quality through review standards, typed contracts and test coverage

HMCTS, Microservice Estate & 3rd Line Support

- 3rd line support across an estate of 70+ services and thousands of repositories; learned Java and Angular on-project and was resolving production incidents in an unfamiliar legacy codebase within three weeks
- Operate ActiveMQ Artemis messaging infrastructure, including a custom management console built against Jolokia/JMX endpoints; currently delivering Angular components for ongoing service modernisation

Product Engineer, RoEx

December 2023 – August 2025

Music AI startup, 8-person team · B2C revenue grew over 200% during tenure

- Researched and prototyped natural-language prompting to extract mixing and mastering preferences, letting users express complex intent without a complex interface, a direct response to user interviews I ran myself

- Rebuilt and A/B tested the core user interface using Microsoft Clarity session data, lifting conversion 20%; rebuilt key components for mobile after analytics showed 60%+ mobile traffic
- Launched client-side buffered audio previews so pro users heard results faster
- Built a deployment pipeline with a Cypress integration suite, cutting critical platform errors by 40% and release cycle time by 25%
- Resolved escalated customer issues at a 95% user-reported satisfaction rating; owned UI/UX and marketing input alongside engineering as team size demanded

Full Stack Software Engineer, PassFort (Moody's Analytics)

August 2022 – December 2023

KYC and compliance platform, team of 10 · Acquired by Moody's Analytics during tenure

- Migrated a legacy configuration system to a DAG-based versioned model in Rust, adding concurrent editing and complete validation, which sharply reduced customer issues around promoting config between environments
- Led full isolation of the largest integration test suite end to end, cutting re-runs from colliding tests by around 30% and running roughly 50% faster through added parallelism
- Practised TDD and held 80% test coverage; monitored services through Grafana, Loki, Datadog and Retool to sustain 99.9%+ uptime
- Ran core team ceremonies, including standups, retros, post-incident reviews and support handover

Junior Full Stack Software Engineer, Netacea

July 2021 – August 2022

AI-powered bot mitigation and cybersecurity, B2B

- Architected and led a from-scratch Akamai EdgeWorkers integration that secured a contract worth over \$1,000,000, profiling it to fit hard limits of 2.5MB memory and 20ms CPU
- Shipped JavaScript edge logic across three CDN platforms (Cloudflare, Akamai and AWS CloudFront) to cut latency and database load; load tested the backend for 100,000+ additional RPS

SELECTED PROJECTS

- **RAG engine** (private): retrieval pipeline making large unfamiliar codebases explorable in natural language
- **GPT from scratch**: transformer implemented from the primitives up, following "Attention Is All You Need"
- **LLM quantisation**: compressing causal LM parameters to cut memory and compute at minimal quality cost
- **Scalar autograd engine**: backpropagation built at the scalar level to internalise the mechanics

SKILLS

Languages TypeScript / JavaScript, Go, Python, Rust, Java, SQL, Bash

AI Engineering RAG and vector retrieval, agent pipelines and tool use, LLM application architecture, open-weight model fine-tuning (Qwen), quantisation, evaluation and deterministic testing of non-deterministic systems, Anthropic and OpenAI APIs, PyTorch, transformer internals, Milvus

Frontend SvelteKit, React, React Native, Angular, Vue, design systems, HTMX

Backend & Data Go, Node, Flask, REST, PostgreSQL, MySQL, MongoDB, Google Datastore

Infrastructure Docker, Kubernetes, Helm, AWS, GCP, GitHub Actions, CircleCI

Observability Grafana, Prometheus, Datadog, Sentry, Loki

Practices TDD, testcontainers, Cypress, Detox, agentic development workflows, Linux

Clearance SC (active)

EDUCATION

BSc Computer Science, University of Reading

2018 – 2021

First Class Honours · Chancellor's Award for the highest weighted grade in Computer Science