Lee Grieve

Applied AI architecture · Enterprise AI deployment · Customer-facing technical leadership

leegrieve.com | hello@leegrieve.com | LinkedIn

Profile

For the last three years I've taken AI from evaluation to production inside The Access Group, a UK software company of about 10,000 people. It's run as 22 semi-autonomous business units, so nothing moves without building consensus across leaders who each have their own priorities, data and risk appetite. My work spans the architecture decisions, the evaluation and guardrails, the governance conversations and the adoption work that decides whether any of it lasts.

Before that, 14 years as a client-facing implementation and escalation consultant, deploying HR and payroll systems into customer estates, cloud and on-premise. Measured results across the automation portfolio so far: 75.47% average effort reduction, £1.43M saved and 46,090 hours released.

Experience

The Access Group | September 2004 to present

Joined when group revenue was around £3m. It now employs about 10,000 people.

AI & Automation Leader | June 2026 to present

Central AI and automation team for Customer Success and Professional Services across 22 business units: 8 engineers in the UK and Kuala Lumpur, reporting to the CCO. Owns the onboarding orchestration platform, the unified customer success workspace, the agent framework and the shared tooling and standards used by delivery pods in each business unit.

Global Head of Technical Operations | December 2023 to June 2026

Led 55+ people across the UK, Kuala Lumpur and other regions while building the automation capability.

Practice Leader, Global Centre of Excellence | January 2020 to December 2023

Led global technical delivery teams across the Access UK product range, including the ANZ enablement programme.

Technical Leader, Centre of Excellence | January 2018 to December 2019

Senior Software Consultant | July 2014 to January 2018

Client-facing implementation of HR and payroll systems into Access Cloud and customer on-premise infrastructure.

Earlier roles | September 2004 to July 2014

HR and Payroll Systems Consultant; Senior Technical Support Consultant; Technical Support Consultant. Owned customer escalations against tight resolution windows.

British Telecom | January 2000 to September 2004

Support Consultant and Site Support Admin. Second-line support across a 500-seat estate and a high-volume contact centre.

Selected work

Architecture and evaluation

Wrote the group's 13-section working reference for its AI platform: the LiteLLM gateway replacing legacy Azure API Management, three access routes (workforce AI, in-product AI, agentic platform) and when to use each, authentication and mandatory metadata patterns, the multi-region model catalogue, and the Kernel Memory retrieval service. Used to onboard engineering teams. I consume and integrate the platform rather than own it.

Argued that defaulting every task to one frontier model, including classification better served by encoder-only or small models, was a cost and capability risk at group scale. Contributed to the multi-model routing that followed; the wider programme cut blended cost per million tokens by 64% in four months, worth about $6M a year.

CSP Navigator. Replaced a manual, email-driven process for mapping products to customer entitlements. The mapping engine is deterministic, config-driven TypeScript because the datasheets are contractual and feed hard-coded entitlements downstream in Service Cloud; the model is confined to conversation and drafting. Next.js and TypeScript, with acceptance fixtures drawn from real product scenarios.

Built a read-only checking sub-agent enforcing nine documented rules, returning PASS, FAIL, BLOCKER or WARNING before any output is finalised.

Co-built the evaluation set for an LLM analytics assistant on Databricks serving seven product families: 25 questions across five difficulty levels with expected SQL patterns and verification criteria, aimed at failures such as ARR double-counted across periods and mishandled negative churn.

Designed an AI categorisation system on Databricks processing c.500,000 records per business unit across 100+ product families, with a four-tier taxonomy that keeps top tiers comparable across the group and lower tiers product-specific.

Took new platform capabilities to the Data Protection Officer as a processor-path and purpose-limitation question before the build decision.

Assessed MCP connector coverage for Professional Services and set Azure DevOps, Databricks and Salesforce as the priority integrations.

Customer-facing and commercial

Took over a live customer escalation at a hospitality account, ran a technical demonstration of the integrations in question and turned the account into a reference advocate.

Tested the go-to-market for a divisional AI rollout and found the blocker was an undefined customer proposition, not product readiness.

Ran a build-versus-buy review of an in-house AI support agent for a care-sector backlog; found it less mature than reported and made the successful case for the commercial product.

Reconciled customer records across three systems with fuzzy matching (79% match rate) to apply delayed price increases to 1,700+ missed customers.

Selected work continued

Adoption and value

Built the AI and automation function from a small squad with no formal remit into a central team of 8 across the UK and Kuala Lumpur, reporting to the Chief Customer Officer.

Delivered 75 automated processes; the largest cut one product's customer onboarding from 40 days to 4.

Built an adoption ladder (from entitled through to sustained use) for board reporting, and used it to separate an instrumentation gap from a genuine adoption problem.

Rolled out a unified customer success workspace (Salesforce, Databricks, Clari) to 350+ users and led the recovery after launch, including an LLM token budget that broke agentic report generation and a disabled Salesforce write-back connector resolved with the CEO.

Enablement

Led the Australia and New Zealand enablement programme, preparing an acquired business for service delivery and sales readiness.

Technical skills

Languages and frameworks: Python (pandas), TypeScript, Next.js, Node/Express

AI systems: LLM APIs, LLM gateways (LiteLLM), model selection and routing, retrieval, MCP servers and connectors, agents and sub-agents, evaluation and guardrail design, Hugging Face, Gradio, local model hosting (Ollama)

Data and cloud: Databricks, Azure (Blob, Front Door, DevOps, API Management), AWS lab environments

Integration and automation: Salesforce, HubSpot, Intercom, n8n, Power Automate

Education

Level 6 AI Engineer apprenticeship, QA Ltd, in progress.

BSc Computer and Information Sciences, Canterbury Christ Church University, 2001 to 2004. Distinction.

Outside work

I run Rusty Spoon Studio, a small indie games studio. Our first game is in public playtest ahead of a Steam release, with Twitch integration that lets viewers vote on what happens next. I've also built and published a KPI positioning agent at kpi.leegrieve.com.