La Fosse

AI Engineering Talent in the UK

What it costs to hire, what it costs to wait, and what to do instead.
H2 2026.

Where this comes from

This briefing is built primarily on three sources that sit inside La Fosse rather than in the public domain. Where market-wide context adds something our own data cannot, such as global technology spend or wider labour market movement, we have drawn on external research too, sourced in full at the end of this document.

Live placement data.

Salary and day rate benchmarks are built from active placement activity across thousands of technology roles on our platform, refreshed twice a year. They reflect what businesses are actually paying, not what they are advertising.

Original research.

In early 2026 we surveyed 2,000 UK tech workers for the AI Reality Check, our research programme on the reality of workplace AI adoption. It sets out the governance gaps, the leadership capability shortfall, and the skills pressure businesses are carrying now rather than the version that appears in vendor marketing.

Practitioner observation.

Our consultants run these processes daily across 22 countries and more than 2,000 companies. Where the data is still forming, and in AI engineering a lot of it is, we say so and tell you what we are seeing instead.

For a private equity backed software business, this last point matters more than usual. 87% of our executive placements sit within PE or VC backed portfolio companies, so the constraints you are working to, the hold period, the board reporting cycle, the pressure to show capability rather than headcount, are the ones we build around by default.

87% **of our executive placements sit within PE or VC backed portfolio companies**

The short version

The salary is not the problem. AI engineering roles in the UK run from £80,000 at mid-level to £150,000 at principal. High, but affordable for any software business already carrying an engineering function.

The cost of acquisition is the problem. Search fees for AI specialists sit at the very top of the market range. A single hire at £120,000 carries £30,000 to £42,000 in fees before the person has written a line of code.

The time is the bigger problem. Four to six months from open requirement to start date is now normal, and that assumes the specification is right first time. Six months of missing capability costs more than the fee does.

The specification is the real problem. Most AI engineer briefs we are asked to work on are describing two or three different jobs at once. That is why processes run long, why offers get declined, and why the role reopens a year later.

Our position: most businesses are trying to buy an AI capability they should be building. The market will sell you one deep specialist at a premium. It will not sell you a team, and it will not sell you one quickly.

> **The salary is not the problem, the cost of acquisition is.**

One market has become two

Generalist technology talent has rarely been more available. Following redundancies through 2025 and a slower correction in hiring, candidate availability is at a multiyear high. For roles where capability is widely distributed this is the strongest buyer's market in several years.

Specialist talent is the opposite story. Salary growth across AI, machine learning and senior data roles has run above 12% in the last twelve months, materially ahead of the wider technology market. Global IT spend is forecast at $6.31 trillion in 2026, up 13.5% year on year, with the sharpest acceleration in AI infrastructure and software. Where investment flows, hiring demand follows.

It is also a smaller pool than the headline numbers suggest. Some of this talent is not choosing between UK employers at all, it is choosing the UK against the US, where the wealth creation opportunity at frontier labs and high-growth companies is pulling experienced AI engineers abroad. That compounds the scarcity described above before a single UK search has even opened.

The practical consequence is that your negotiating position is completely different depending on which side of that line the role sits. Treating both with the same hiring process, the same timeline and the same offer strategy is where most businesses lose the specialist candidates.

What AI Engineers cost: **permanent**

Low to high reflects level and scope, not negotiation room.

Role	Low	High
Chief Data Officer	£180,000	£300,000
Head of AI	£120,000	£200,000
Head of Machine Learning	£120,000	£200,000
AI Architect	£110,000	£150,000
Machine Learning Engineer	£90,000	£150,000
AI Engineer, mid to principal	£80,000	£150,000
ML Researcher	£90,000	£140,000
Computer Vision Engineer	£90,000	£130,000
AI Product Manager	£90,000	£130,000
MLOps Engineer	£80,000	£120,000
AI Governance Lead	£90,000	£110,000
NLP Engineer	£80,000	£100,000
AI Business Analyst	£60,000	£90,000
Data Engineer	£50,000	£90,000
Analytics Engineer	£50,000	£85,000

AI hiring remains heavily concentrated in London and the South East, so the regional discount available in other engineering disciplines is thin here. Remote-first hiring widens the pool but does not lower the price, because everyone else is hiring remotely for the same people.

This table reflects the broad UK market. At the top of it, the numbers break down. AI labs and frontier companies, banking and hedge funds, and the highest-growth start-ups routinely pay £100,000 to £200,000 or more for the same seniority band, because they are not competing against the UK market, they are competing against the US one. If your search is targeting that calibre of candidate, budget against that number, not the benchmark above it.

What AI Engineers cost: **on contract**

Low to high reflects level and scope, not negotiation room.

Role	Inside IR35	Outside IR35
Head of AI	£850 to £1,000	£800 to £950
AI Architect	£850 to £1,100	£750 to £1,000
AI Engineer, mid to principal	£800 to £900	£700 to £800
Machine Learning Engineer	£800 to £900	£700 to £800
MLOps Engineer	£750 to £850	£700 to £800
AI Governance Lead	£750 to £850	£650 to £750
NLP Engineer	£700 to £800	£600 to £700
AI Business Analyst	£650 to £750	£550 to £650

Contract is the fastest route to capability and the most expensive way to hold it. It works for finite work with a clear end point. It does not work as a permanent answer to a permanent problem, which is how most businesses end up using it.

What the title actually means

AI Engineer is not yet a standardised role, which is why the band is so wide and why benchmarking against a job title alone will mislead you. We are also seeing the title used loosely on the client side. Automation roles, support engineering roles and data scientist roles are regularly being labelled AI Engineer in job specs, which inflates the apparent talent pool and sets a false benchmark before a search has even started. Here is how the levels separate in practice.

Mid-level, £80,000 to £100,000. Builds on existing models and APIs. Retrieval systems, prompt and evaluation harnesses, integration into product surfaces. Ships features. Does not set architecture.

Senior, £100,000 to £125,000. Owns a system end to end. Decides between prompting, retrieval and fine-tuning on cost, latency and accuracy grounds. Builds the evaluation that tells you whether it works. Works directly with product on what is feasible.

Principal, £125,000 to £150,000. Sets technical direction for AI across the product. Makes the buy versus build calls on models and infrastructure. Accountable for cost per inference and for what happens when the model is confidently wrong.

These are different from the roles they are routinely confused with. A Machine Learning Engineer trains and deploys models and owns the pipeline, closer to data engineering than to product. An MLOps Engineer owns infrastructure, monitoring, drift and cost. A Data Scientist runs analysis and experimentation and is rarely building production systems. An AI Product Manager decides what to build and defines what an acceptable failure looks like. An AI Governance Lead owns policy, oversight, documentation and regulatory readiness.

The most common brief we are asked to fill describes two or three of these in one job. It fails for a predictable reason. Either you find nobody, or you find a generalist who is credible at all of it and excellent at none of it. Splitting the brief before you go to market is the single highest-return hour anyone spends on an AI hire.

What the salary number does not tell you

Fees. Contingent recruitment across the wider technology market typically runs between 15% and 25%. Retained and specialist search sits between 25% and 33%. Elsewhere in the market, AI engineering is transacting at the very top of that range, and 30% or more is increasingly being quoted and paid on search-led processes. That is not opportunism. It reflects genuine supply constraint at the senior end.

We do not deviate from our standard fee structure for AI engineering hires with existing clients. That means 18% to 25%, the same as any other technology search, not a premium for the difficulty of the role. Businesses paying up towards 30% or more elsewhere are not paying for better access to this market. They are paying for the same access we already have, priced the way we would price any other search.

Time. Four to six months to start date is normal, before notice periods of one to three months on top. Model that as capability arriving in month five to seven, not month one.

Failed process risk. Candidates in this market typically run three or more live processes. Counter-offers are common and increasingly effective. A process that collapses at offer stage costs the elapsed time twice.

Employer cost above base. National Insurance, pension, equity and the re-benchmark that arrives twelve months later when the market has moved again. Budget the base, then add roughly a quarter before you have a true number.

18-25%

is our standard fee structure for AI engineering hires with existing clients.

The comparison that actually matters

There is no single right answer to a capability gap. **Hire** buys you a permanent specialist at a premium price and a slow start date. **Borrow** buys you the same depth on contract, faster, but you lose the capability the day the contract ends. **Build** grows the capability internally, slower to reach full depth but retained for good and cheaper by a wide margin. Most businesses default to hire for everything. The organisations getting AI capability right choose deliberately between the three, role by role.

One mid-level AI engineering capability, held for twelve months, three ways. Illustrative, using the benchmarks above and 220 billable days.

	Hire	Borrow	Build
Model	Permanent hire	Contractor	Trained associate
Base or rate	£120,000	£800 per day	£345 per day
Fee or margin	£36,000 at 30%	In the rate	In the rate
Twelve-month cost	£156,000	£176,000	£75,900
Live from	Month five to seven	Week two to four	Week two to four
Capability retained	Yes	No	Yes
Depth	High	High	Building

These are not interchangeable, and the table should not be read as saying build is three ways better. A trained associate does not replace a principal engineer, and no amount of programme design pretends otherwise. What the numbers show is that the deep specialist is the most expensive and slowest thing you can buy, so you should buy as few of them as the problem genuinely requires.

For every capability gap, three questions decide the model. Will we need this capability in three years? Is the talent available at reasonable market cost? Is the requirement finite or ongoing? Finite and specialist means borrow. Ongoing and specialist means hire, and start now. Ongoing and learnable means build. Most workforce planning failures happen because these questions are never asked out loud.

Most AI hires are the wrong shape

The common failure is not paying too little. It is hiring one expensive specialist into an organisation where nobody can brief them, evaluate their work, or take what they build into production. The specialist becomes a bottleneck, then becomes bored, then leaves. The role reopens twelve months later at a higher price.

The organisations getting this right design in capability layers. A small number of deep specialists, hired properly and paid at market. A much broader layer of technically fluent people around them who can hold the AI conversation without being AI engineers. That second layer is buildable through upskilling and structured early-career routes, and it is where the retention argument sits too. People stay where investment in progression is visible.

The roles most businesses are under-hiring for are not the engineering ones. AI Governance Lead, AI Product Manager and the emerging orchestration and human-in-the-loop roles sit between £90,000 and £130,000, are easier to fill than an AI engineer, and unblock the specialists you already have.

Adoption is a metric. Capability is a strategy. Most businesses are still reporting the first and hoping it reads as the second.

The organisations getting this right design in capability layers.

How to assess an AI engineer when nobody internally can

This is the question we get asked privately more than any other, usually after a first hire has not worked out. If your existing engineering leadership has not built with these systems, they cannot reliably grade the answers, and strong interviewers default to testing what they know instead.

Do not test model architecture unless you are training models. Almost nobody outside a research lab is, and the candidates who perform best on those questions are frequently the ones least interested in shipping product.

- **Test evaluation thinking.** How would you know this works? How would you know when it stops working? Weak candidates describe the build. Strong ones describe the measurement.
- **Test failure handling.** What happens when the model is confidently wrong in front of a customer? Look for someone who designs for that rather than someone who promises it will not happen.
- **Test cost and latency awareness.** Inference cost, caching, when a smaller model is the right answer, and when the right answer is not to use a model at all.
- **Test the handover.** Can they explain a technical trade-off to a commercial stakeholder without either patronising them or hiding behind the jargon? This is the fastest-growing part of the specification and the hardest to teach.

Where the internal bench genuinely cannot assess, bring in an external technical assessor for the second stage. It costs a fraction of a failed hire and it protects the offer stage, because candidates in this market can tell within twenty minutes whether the people interviewing them understand the work.

Hiring them is half the job

Retention pressure on senior technology hires is rising, and replacement cost is rising with it. Compensation still matters, but progression, recognition, flexibility and visible investment in development are increasingly what decides whether someone stays.

For AI engineers specifically, a few levers move the needle disproportionately, and most of them are cheap.

- **Access.** A real budget for computer and model access. Being told to build with AI and then being rate-limited is the most reliable way to lose one of these hires.
- **Exposure to the cutting edge.** Partnerships with the likes of OpenAI and Anthropic, access to frontier models and tooling, and visible proximity to where the technology is actually moving. This population is choosing employers partly on what they get to build with, not only on what they get paid.
- **Autonomy over the stack, within commercial guardrails.** These are people who left somewhere else because the decision was made above them.
- **Proximity to the decision.** If the AI roadmap is set two floors up and handed down, the person you paid a premium for is an implementer, and they will price themselves accordingly at their next move.
- **Wealth creation, not just compensation.** Long-term incentive plans, equity and genuine ownership matter disproportionately to this group, and increasingly compete directly with what US employers are offering. A strong base salary alone is no longer enough to hold the best of this population
- **Visible external profile.** Conference slots, writing, community. In a PE-backed business this is often easier to give than equity, and for this population it is frequently worth more.

Assume the counter-offer. Build the offer to survive one, and structure a re-benchmark conversation at twelve months before someone else has it for you.

The supply lever most businesses ignore

Women remain materially underrepresented across AI and machine learning disciplines, and in our research only 28% of women in technology see a clear path to the C-suite. Mid-career is where the attrition concentrates, which is exactly the level at which AI engineering demand sits.

Read commercially rather than as a values point, that is a supply problem. If you hire from the standard pool you are competing for the same shortlist as everyone else, at the top of the fee range, on a six-month clock. Businesses widening the pool are hiring faster and paying less, and they are doing it by fixing progression and sponsorship rather than by running another entry-level campaign.

The same logic applies to career changers and structured early-career routes. Apprenticeships and associate programmes are increasingly being used as a strategic response to long-term specialist shortage, not as a diversity initiative that sits alongside the real hiring plan.

Why some of this talent leaves the UK altogether

Part of the supply problem is not competition between UK employers, it is competition with the US. The scale of the volume and wealth creation opportunity available to AI engineers in the US, through equity in frontier labs and high-growth companies, is drawing UK talent abroad, and that is compounding the domestic shortage described throughout this briefing.

This is not something any single business can fix on its own, but it is worth understanding as context for why the search is as hard as it is. Closing the gap needs movement in three places: continued government investment in AI infrastructure, innovation and UK scale-ups; more growth capital so AI companies can scale here rather than relocating to raise it; and clearer career pathways from the UK's leading universities, UCL, Imperial, Oxford and Cambridge among them, directly into AI roles rather than into companies that later move talent overseas.

The regulatory clock

The EU AI Act applies to AI systems used across recruitment, candidate screening, performance monitoring, promotion decisions and workforce analytics. Businesses operating internationally or hiring into EU markets fall within scope regardless of where the parent sits. The UK approach differs, but workforce capability and AI governance are becoming strategic priorities across both public and private sectors either way.

For a business with European operations this is the least discretionary item on the list, and it is also the cheapest to get ahead of. Over the next ninety days: audit the AI-enabled HR and hiring tools already running, document human oversight for each, build AI literacy across HR and leadership, review candidate transparency, and assign a named owner for ongoing compliance. None of that requires a hire. All of it requires someone to be accountable.

What to do in the next 90 days

1. Split the requirement

Write down which parts of the problem genuinely need a deep specialist and which need a technically fluent person with good judgement. Most briefs shrink by a third at this step.

2. Run the three questions against every gap

Then decide hire, borrow or build role by role rather than as a policy.

3. Secure the one or two hires that genuinely have to be hires

Candidate availability is still elevated across the wider market, but the window for specialist talent is narrowing as confidence returns.

4. Fix the assessment process before you go to market

This includes external technical assessment if the internal bench cannot grade the answers.

5. Audit the AI tools already running

Across hiring, HR and performance monitoring, and name an owner for oversight.

6. Treat AI literacy as a capability expectation for the existing team

Rather than a training initiative for next year. The premium is increasingly paid for people who combine technical fluency with commercial judgement, and you already employ some of them.

Working with La Fosse

We are built to work across all three models in the comparison table, which matters because almost no AI capability problem is solved by one of them alone.

La Fosse Executive places leadership and interim talent, with 150+ specialist consultants and 15+ years of experience across 22 countries and more than 2,000 companies. 87% of those placements sit in PE or VC backed portfolios.

La Fosse Recruitment covers the specialist and senior hires, from AI and machine learning engineering through to governance, architecture and product.

Inovus by La Fosse builds and manages delivery teams for finite, transformation-led work, and hands the capability back rather than keeping the dependency.

Academy by La Fosse builds the layer underneath. Trained associates embedded into client teams, upskilling for existing engineers, and apprenticeships that turn levy funding into structured development. Associates and apprentices work across data, technology and operations for organisations including UCL, Royal London and Leidos, and the cohorts are deliberately built wide: 46% career changers, 38% women, drawn substantially from backgrounds the standard pipeline does not reach.

Alongside the commercial work we run the AI Reality Check research programme, the AI Sense Check benchmarking tool, and UNBOUND, our community for women in technology and transformation. The next phase of the research is Beyond Headcount, a blueprint built with senior leaders from UK enterprise, private equity and the public sector on how the workforce needs to reshape for 2027.

The most useful next step is usually the smallest one. Send us the requirement as it currently stands and we will tell you which parts of it are hire, which are borrow, which are build, and what each will realistically cost and take. That conversation is free and it is worth having before the fee conversation, not after it.

Notes on data

Salary and day rate ranges are drawn from the La Fosse H2 2026 Workforce Insights Report, combining proprietary placement and hiring data across thousands of roles with consultant observation from live market conversations and external labour market reporting, including Gartner, KPMG and REC, McKinsey, PwC, LinkedIn, DSIT and BCS. Fee ranges reflect observed market practice rather than any single firm's terms. The twelve-month cost comparison is illustrative.

One caveat worth stating plainly. Benchmarks for the newest roles, such as Agentic Operations Manager and Human-in-the-Loop Designer, are still forming, because the roles are younger than the data. The demand is not still forming.

La Fosse