



GLOBAL AI TALENT SURVEY

Insights from numi's network Surveys

INTRODUCTION

Over 500 of the world's top product, engineering, and founder-level executives across locations, including the GCC, Berlin, New York, London, and Amsterdam, were surveyed.

Our survey drew exclusively from executives at venture-backed or private equity-funded companies, all at Series A or beyond, with credible ARR and teams ranging from 50 to 1,000+.

The results provide a high-resolution view into how the world's most ambitious technology leaders are building AI-native teams and products.




AI HIRING & TEAM STRUCTURE

THE RISE OF THE APPLIED AI ENGINEER

Machine Learning Engineers continue to be the workhorse of AI hiring (48%), with sharp demand also building for Data/AI Platform Architects (33%) and AI Product Managers (31%) – signalling a transition from experimentation to productisation. Product-Focused ML Engineers (24%) round out the shift. The fact that 21% are not hiring AI roles is equally revealing: a clear divide is forming between the builders and the bystanders. Teams that haven't made foundational AI hires by 2025 risk being permanently left behind.

MID-SENIOR TALENT IS THE CORE OF EXECUTION

Over half of all AI hires are in the 3-6 year experience bracket. Senior hires (7-12 years) are not far behind at 43%. Entry-level hiring sits at only 18%, indicating that AI isn't a playground for generalists fresh out of university. Companies want velocity, not training projects.



AI IS NO LONGER A SIDECAR



78%

Are embedded into core product/engineering teams

75% embed AI talent directly within core product and engineering teams.

AI is a layer within every meaningful product. Meanwhile, 36% still place AI within R&D labs, suggesting a bifurcation: innovation hubs versus embedded execution. The hybrid team model (31%) is winning over centralised AI centres (16%) as companies move toward decentralised deployment with centralised oversight.

52%

Are embedded into Data and analytics teams

36%

Are embedded into Innovation or R&D labs






HIRING TRENDS

Machine Learning Engineer and AI Product Manager are highly sought-after AI roles for hiring or planning to hire in 2025.

Hiring for AI roles is concentrated at the Mid (3–6 years) and Senior (7–12 years) experience levels.

AI talent is primarily being embedded within core product/engineering teams and Data and analytics teams.

1 in 5 responders are not planning on hiring AI talent this year.



TOOLS, TECHNOLOGY &
INFRASTRUCTURE

OPENAI IS THE DEFAULT, BUT COMPLEXITY IS RISING

78% of respondents use OpenAI's APIs. But beneath that headline is fragmentation: 48% use Anthropic, 31% are already deploying RAG pipelines, 26% use vector DBs, and 26% are experimenting with LangChain.

Hugging Face and autonomous agents lag behind. The clear implication? The LLM gold rush is well underway, but toolchain maturity is still in flux.

CLOUD DOMINATES – FOR NOW

Public cloud remains the backbone of AI infrastructure (70%). Serverless and fully managed platforms are emerging, but their adoption is still in early innings. Only 8% are on-prem. The AI infra wars are just getting started – and consolidation is coming.

FINE-TUNING IS THE NEW FRONTIER

A third of companies are blending foundational models, fine-tuning, and custom training. Only 5% are training from scratch. This is the new arms race: the ability to selectively fine-tune open-source models for unique data moats. That 24% are already doing this marks a tipping point.



 OpenAI 75.2%

 Claude 46.8%

   Azure 72%

TALENT STRATEGY, SOURCING & BRANDING



THE NETWORK EFFECT OF GREAT TALENT

71% cite personal referrals as their best hiring channel. AI talent moves through high-trust, reputation-based networks. It's not enough to post a job – you need to be in the room, or know someone who is.

LinkedIn is helpful (39%), but far from dominant. Executive search firms (18%) and university partnerships (14%) trail, revealing that the old pipelines don't match AI's velocity.



39%



18%



14%

REMOTE-FIRST VS. DEEP ECOSYSTEMS

London and San Francisco remain dominant AI hubs, but 40% of respondents said they hire remote-first globally. The takeaway? You either bet on density (London, Bay Area) or you win by accessing global edge talent. Dubai, Singapore, and Bangalore are on the rise but not yet top-tier.



45% of people perceived London / Cambridge as having the highest concentration of elite AI talent




synthesia



37% of people perceived San Francisco / Bay Area as having the highest concentration of elite AI talent

TEMPUS

Leena AI



WHY TALENT SAYS YES

It's not just about comp: 60% of AI professionals are motivated by mission and purpose, 50% by product impact, and 39% by visionary technical leadership.

Compensation matters (35%), but it's not the primary lever. The best people want to work where their fingerprints matter.

THE HIRING BOTTLENECKS ARE INTERNAL

The biggest blockers to hiring aren't FAANG or GPU shortages. They're internal: 30% cite lack of AI maturity, 27% point to comp benchmarks, and 14% to brand reputation.

Companies that don't have a cohesive AI story, or can't articulate their vision, simply lose.



EVALUATION & INTERVIEWING

THE BAR IS HIGH AND MULTI-DIMENSIONAL

Technical assessments (coding, system design, use case builds) are standard. But product thinking, collaboration, and ethical reasoning are now baseline expectations. 51% assess collaboration and product impact, 45% evaluate ML architecture fluency, and 53% run algorithmic interviews. Elite hiring means assessing for breadth and execution.

NON-NEGOTIABLES FOR 2025 AI TALENT

The new table stakes: awareness of hallucination and safety risks (51%), prompt engineering (49%), fine-tuning LLMs (46%), deploying production-ready models (41%), and API engineering (39%). This isn't about hobbyist hackers. It's about end-to-end builders.

PRODUCT+AI IS THE WINNING FORMULA

52.5% now prioritise product engineers fluent in AI over deep specialists. Layered teams still exist (21%), but it's clear: the companies moving fastest are those embedding AI fluency across the stack.



COMPENSATION & MARKET TRENDS

COMPENSATION AND INCENTIVES

The average total compensation for Entry-level AI/ML Engineers is most frequently reported as <\$100k.

For Senior AI Engineers and AI Product Managers, the \$100k–150k range is the most common average total compensation.


AI compensation expectations have increased moderately over the past 12 months.

Most organisations do not currently offer AI-specific incentives and do not intend to.



COMPENSATION AND INCENTIVES

The below charts give a benchmark for base salary + equity on average from our respondents' companies for the following in-demand roles:



FUTURE EVOLUTION

AI teams are expected to evolve towards being embedded in every product squad, and a rise of product engineers with AI fluency. Investment in Autonomous agents & copilots and LLM agent developers is a priority for future AI talent development. Fintech & Insurtech, SaaS & Productivity, and Cybersecurity are the verticals expected to hire the most AI talent in the next 24 months.

7 IN 10 PEOPLE BELIEVE AI WILL BE EMBEDDED IN EVERY PRODUCT SQUAD




7 IN 10 PEOPLE THINK THERE WILL BE A RISE OF PRODUCT ENGINEERS WITH AI FLUENCY



60% OF PEOPLE SEE AI COPILOTS AUGMENTING INTERNAL TOOLS



32% OF RESPONDERS SEE A GROWTH OF CROSS-FUNCTIONAL AGENTS/AGENT TEAMS





SKILLS & EVALUATION

Awareness of hallucination, bias, and safety risks, Prompt engineering, and LLM fine-tuning & evaluation are considered non-negotiable skills for applied AI roles.

There is a trend towards prioritising product engineers with AI fluency or valuing a hybrid product+AI thinking skillset.

Fintech & Insurtech leads by a wide margin, with SaaS & Productivity and Cybersecurity rounding out the top three – confirming AI hiring concentrates where the cost of getting it wrong is highest and the pressure to ship it well is greatest.



WANT TO BUILD LIKE OpenAI?

HIRE TALENT FROM:



They look for core technical skills, including:
Python · Java · C++ · Algorithms · Data Analysis · ML · JavaScript

THEY'VE STUDIED AT:

UC Berkeley MIT Stanford Harvard Carnegie Mellon

Top locations to source AI talent from are San Francisco, London, and New York

KEY HIRES YOU NEED TO MAKE

AI Engineers Software Engineers Research Engineers
Product Managers Members of Technical Staff

"We partnered with numi because they have a strong network and deep knowledge of the locations where we hire.

They quickly understood the type of experienced engineers we were looking for and quickly found great candidates, which led to successful hires. They were always very responsive, and they understood the profiles we were from a technical perspective very well. Working with them was easy and enjoyable."

Sara Vošinek Gašpar, Head of Talent



THE FUTURE BELONGS TO AI-NATIVE COMPANIES

This isn't an era of optional AI adoption. It's a re-architecture of company DNA.

The data makes it clear: companies winning in AI are not just hiring engineers. They're redesigning their org charts, evolving their product philosophies, and building for compounding capability.

From London to San Francisco to Dubai, there is a visible separation between companies embedding AI into their bloodstream versus those still flirting with pilot projects. Founders and investors should be asking one thing: Are we building an AI-native company or retrofitting yesterday's model?

AI hiring is no longer experimental. It's existential.



HELLO@WEARENUMI.COM

