

HIRING AI TALENT IN 2026

USA MARKET REPORT



TABLE OF CONTENTS

INTRODUCTION	1
EXECUTIVE SUMMARY	2
THE USA AI MARKET	3
HIRING ADVICE	4
ROLE DESIGN	5
SALARY BENCHMARKS	6
TRENDS	7
BLOCKERS TO SUCCESSFUL HIRING	8
FINAL THOUGHTS	9

INTRODUCTION



DAVE CARROLL
Senior Consultant
KDR Talent Solutions

This report examines the US AI talent market across the main hubs of San Francisco Bay Area, New York, Seattle, Boston, Austin, Dallas–Fort Worth, Houston, Atlanta, Denver/Boulder, Raleigh–Durham, Chicago, Los Angeles and Washington DC/Northern Virginia. Against a backdrop of the wider country and the competition this creates.

The focus is on AI Engineering, Machine Learning Engineering, LLM Engineering, MLOps, Computer Vision and AI leadership roles.

Artificial intelligence has moved from experimentation into active workforce planning. Employers are now building AI into products, platforms, workflows, infrastructure and operational systems.

That shift has changed the talent market. Companies increasingly need people who can design, deploy, monitor and maintain AI systems in production — not just experiment with models.

Nearly 125,000 active US job postings required AI skills in May 2025, a clear signal that AI capability has become a mainstream hiring priority and not just a specialist technology function.

EXECUTIVE SUMMARY

The US AI talent market is really a collection of regional ecosystems, each with different strengths, salary pressures and candidate expectations.

The San Francisco Bay Area for example remains the centre of gravity for frontier AI, research labs, model companies and AI infrastructure whereas Seattle is especially strong for cloud, MLOps and ML platform talent. New York is highly relevant for applied AI in finance, legal, media and enterprise software and Boston remains strong for research-led AI, robotics, biotech, healthcare and deep tech.

Applied AI demand is also spreading quickly. Austin is a major market for product-led AI, SaaS, robotics and startups. Dallas–Fort Worth offers scale and enterprise technology growth.

For hiring teams, the implication is clear: employers are competing nationally. Salary matters, but candidates also assess the technical challenge, AI roadmap, leadership, flexibility and whether the work is genuinely production-focused.



WHAT THIS MEANS FOR HIRING TEAMS

The strongest AI candidates are rarely only comparing roles within one city. They compare opportunities by technical challenge, compensation, flexibility, leadership quality and whether the AI work is real.

For hiring managers, the candidate message must be specific. Strong candidates want to know what problem is being solved, what has already been built, whether the system is in production, what data is available, who owns the roadmap, how success is measured and what ownership the role carries.

The companies that win AI talent can explain the problem clearly, move quickly, benchmark compensation realistically and show that the work is meaningful, funded and technically credible.

*Quick
Tips*

Benchmark compensation against the right market

AI candidates often compare opportunities nationally. A company in Dallas, Houston, Atlanta or Raleigh may benchmark locally, while the candidate compares the role with remote opportunities from New York, Seattle, San Francisco or venture-backed AI startups.

The highest salary does not always win. If compensation is below market, the role needs to offer something stronger elsewhere: equity, mission, flexibility, ownership, technical challenge or progression. Employers should benchmark against the actual talent pool, not simply their nearest city.

5 AREAS TO FOCUS ON FOR HIRING MANAGERS

We would suggest there are five areas hiring managers should focus on.



Be clear about the problem, not just the job title.

AI job titles are still inconsistent. “AI Engineer” can mean product engineering, machine learning, LLM integration, model deployment or platform work. “GenAI Engineer” can mean anything from basic API integration to complex RAG architecture and evaluation.

Candidates will respond better when the hiring manager can clearly explain the problem they are trying to solve.

In the past companies may have positioned roles as slightly more generic, for example:

“We are hiring an AI Engineer to help us with GenAI.”

Now stronger messages win such as: “We are building an LLM-powered workflow tool for our customers. The first challenge is improving retrieval quality, reducing hallucinations and creating an evaluation framework before we scale it across the product.”

That level of clarity gives the candidate confidence that the company understands the role.

Show that the AI roadmap is real

Senior candidates want to know whether AI is genuinely important to the business or simply a board-level buzzword. They will look for signs that the company has executive support, usable data, engineering investment and a clear reason for building AI.

Hiring managers should be ready to explain things like what has already been built v’s what is still experimental. What will move into production, who owns the AI strategy. What budget and infrastructure are in place and how the role connects to product, revenue, efficiency or customer value.

This does not mean the company needs to have everything solved. In fact, many candidates enjoy joining at an early stage. But they need to believe the opportunity is real.

5 AREAS TO FOCUS ON FOR HIRING MANAGERS

Sell the technical challenge

The best AI candidates are motivated by difficult, meaningful problems. They want to know what makes the work interesting.

The hiring message should be tailored to the candidate profile. A Computer Vision Engineer may care deeply about the physical-world use case whereas an MLOps Engineer may care more about platform maturity, tooling and ownership for example. It may seem obvious but many job adverts don't play to the candidate's motivations and preferences strongly enough, making it hard to stand out.



Move quickly and keep the process relevant

Strong AI candidates are often in multiple conversations. A slow or generic hiring process creates risk. The best processes are structured but not excessive. They test the right things and give the candidate confidence in the team.

Some example processes that have worked in the past have tailored the structure for the role. For an LLM Engineer, the process could explore RAG, evaluation, model reliability, data privacy and production trade-offs. For an AI leader, it should explore roadmap, team design, governance, stakeholder influence and commercial outcomes.

A poor process asks generic coding questions that do not reflect the job. A strong process helps both sides understand whether the candidate can solve the actual problem.

In AI hiring, the companies that win talent are not always the companies with the biggest brand. They are the companies that can explain the problem clearly, move quickly, benchmark compensation realistically and show candidates that the AI work is real.



ROLE DESIGN

Before going to market, hiring managers should define the role around the business problem rather than the job title. In AI hiring, titles are still inconsistent, so starting with “we need an AI Engineer” or “we need an LLM Engineer” can lead to a vague brief and a difficult search. A better approach is to work through five questions.

1. Define the business problem

Start by identifying what the person needs to help the business achieve.

2. Identify the centre of gravity

Once the problem is clear, decide where the role naturally sits. If the priority is roadmap, governance, team structure and commercial direction, it may sit in AI Leadership. This helps avoid building an unrealistic “everything AI” job description.

3. Separate essential skills from useful extras

AI briefs often become overloaded because every stakeholder adds a requirement. Hiring managers should separate what the person must have from what would simply be helpful.

For example, a role may genuinely require RAG, Python, cloud deployment and model evaluation. It may not also need computer vision, fine-tuning, Kubernetes, data engineering, product management and team leadership.

The more requirements added, the smaller and more expensive the candidate pool becomes.

4. Match the interview process to the actual work

The assessment should test the problem the candidate will solve. A generic coding test may not be useful if the role is about LLM reliability, MLOps architecture or computer vision deployment, for example. A relevant process improves candidate experience and helps the hiring team make better decisions.

5. Benchmark the role against the right talent pool

The final step is compensation and positioning. Hiring managers should benchmark the role against the actual level of specialism required, not against a generic software, data or local salary range. The more specialist and production-critical the role, the more important it is to benchmark against the right market.

The practical test

Before opening an AI role, hiring managers should be able to answer this in one sentence:

“We are hiring this person to solve [specific problem] by doing [core work], and success in the first six to twelve months looks like [measurable outcome].” For example: “We are hiring an LLM Engineer to improve retrieval quality, evaluation and reliability for our customer-facing AI assistant, with success measured by improved answer accuracy, lower escalation rates and production readiness within six months.” That level of clarity makes the role easier to market, easier to assess and easier to fill.



NEED ACCURATE SALARY BENCHMARKING?

We provide bespoke compensation reports based on your role, location, industry, and technical requirements, helping you benchmark against the right talent pool rather than broad national averages.

AI compensation varies significantly based on factors such as role scope, production experience, location, industry, company type, and the technologies required. As a result, broad national salary ranges often fail to reflect the realities of individual hiring markets.

KDR provides bespoke compensation benchmarking tailored to your specific hiring requirements, including:

- Role and level
- Location (local or national market)
- Industry
- Company type (startup, consultancy, enterprise, etc.)
- Required technical skills (e.g. Databricks, MLOps, LLMs)
- Current market competition

This enables hiring teams to benchmark against the talent they are actually competing for, rather than relying on broad national averages.

WHAT ARE THE TRENDS IN HIRING AFFECTING THIS MARKET?



TREND 1: AI SKILLS ARE NOW A MAINSTREAM HIRING REQUIREMENT

AI hiring is no longer limited to frontier labs or pure AI startups. Employers across software, finance, healthcare, energy, logistics, manufacturing, cybersecurity and enterprise technology are assessing how AI fits into their workforce plans.

Competition is broader than many expect. A company in Atlanta, Austin or Chicago may be competing with local employers, national remote teams, Big Tech and venture-backed AI-native businesses.

TREND 2: PRODUCTION EXPERIENCE MATTERS

AI/ML employers increasingly need candidates who can support production systems. Interview processes should test what candidates have shipped, how they handled deployment, how they measured performance and how they maintained reliability after launch.

TREND 3: LLM HIRING REQUIRES CAREFUL SCOPING

LLM Engineer remains less standardised than Machine Learning Engineer. Basic GenAI integration should not be scoped like a role requiring RAG architecture, fine-tuning, evaluation systems, routing, guardrails or LLM infrastructure. The more advanced the requirement, the more likely employers are to face national salary competition.

MLOps is also becoming business-critical as deployment, monitoring, model lifecycle, governance and platform reliability grow in importance. Strong candidates increasingly expect a real use case, usable data, technical leadership and a credible roadmap.

WHAT ARE THE MAIN BLOCKERS TO SUCCESSFUL HIRING



HIRE FOR POTENTIAL, NOT PERFECTION

One of the biggest blockers in AI hiring is the temptation to search for the perfect candidate. AI job descriptions often combine LLMs, machine learning, MLOps, cloud architecture, data engineering, software engineering, product thinking, domain expertise and leadership.

Very few candidates match every requirement, and those who do are expensive, highly sought after and usually already in strong roles.

Hiring only for perfect experience slows the process, reduces the candidate pool and risks overlooking high-potential talent who could grow quickly. This is especially true at junior and mid-level, where practical ability, learning speed and strong engineering fundamentals can matter more than exact tool experience.

Hiring managers should separate must-have capability from trainable experience.

Role ambiguity is another major blocker. Employers may know they need AI talent without defining whether the work centres on product AI, LLMs, machine learning, MLOps, computer vision or leadership. The result is vague job descriptions, misaligned interviews and unrealistic salary expectations.

WHAT ARE THE MAIN BLOCKERS TO SUCCESSFUL HIRING



National competition is a further blocker. Strong AI candidates compare roles against remote-first companies, Big Tech, AI-native startups and higher-paying US hubs, especially in LLM Engineering, MLOps, Staff-level ML and Computer Vision.

There is also a shortage of production AI experience. Many candidates have AI interest or prototype experience, but fewer have deployed, monitored and scaled systems in real environments. Production GenAI, RAG, evaluation, MLOps and AI infrastructure talent therefore remains highly competitive.

Overloaded job descriptions narrow the candidate pool and push compensation higher. Slow interview processes, unclear AI roadmaps and compensation misalignment also lose candidates.

Strong candidates want to understand the business problem, technical challenge, data environment and what success looks like. For hiring managers, the message is simple: define the role clearly, benchmark against national AI talent, move quickly and show that the work is real, funded and technically credible.

FINAL THOUGHTS



The US AI hiring opportunity is significant, but so is the competition. AI talent is national, salary expectations vary widely, and the strongest candidates expect clarity about the role, technical challenge, data environment, roadmap and compensation.

The Bay Area, New York, Seattle and Boston create the strongest salary pressure. Austin is a clear applied AI growth market. Dallas-Fort Worth, Houston, Atlanta, Denver/Boulder, Raleigh-Durham, Chicago, Los Angeles and Washington DC/Northern Virginia each offer distinct opportunities depending on the role and sector.

Employers that define roles accurately, benchmark compensation realistically and move decisively will be best placed to turn AI ambition into delivery.

KDR helps clients define roles, benchmark compensation, understand candidate expectations, access specialist talent and reduce time-to-hire.

About Dave Carroll

Dave Carroll leads KDR Talent Solutions' US AI recruitment practice, supporting companies across the US as they build AI capability across engineering, platform and leadership functions.

He advises on market conditions, role definition, compensation benchmarking and access to specialist AI talent.

Contact:

Dave.Carroll@kdrtalentsolutions.com

Senior Consultant

Let's connect on LinkedIn

