

karat[^]

AI Workforce Transformation Report

2026-2027



Executive Summary

The first phase of enterprise AI was dominated by experimentation. The next phase will be about economics.

Karat's 2026-27 AI Workforce Transformation Survey, conducted by The Harris Poll, asked 157 technology leaders across the U.S. and China what increased AI usage and productivity mean for the design, cost, and value of engineering organizations.

We also sat down with several prominent engineering and talent leaders across the industry to discuss the findings and provide commentary.

Three findings stand out.

1. AI is creating capacity, not eliminating work

In the U.S., 90% of leaders are using AI to do more while maintaining or expanding their team size. Only 10% are maintaining their current output while reducing headcounts. What's more, engineering work isn't disappearing. Engineers are spending less time writing code, but more time reviewing AI outputs, refining requirements, testing, and conducting security reviews.

2. AI is stretching the engineering performance curve

Globally, 75% now say a strong engineer is worth at least 3x total compensation. In the U.S., the share saying a strong engineer is worth at least 5x has jumped from 18% in 2025 to 29% in 2026. Organizations are responding by growing senior IC headcounts while entry-level roles face pressure. The premium on talent quality also extends beyond the FTE workforce: 54% of CIOs are very satisfied with the quality of their full-time engineers, compared with just 37% for contractors.

3. The total costs of Human + AI engineering are changing faster than the ability to benchmark and measure the impact

CIOs are increasingly managing AI as part of the total cost of engineering, but most lack clear benchmarks for what "good" looks like. Nearly three-quarters believe their organizations are ahead of their industry in leveraging AI, while just 5% believe they are behind. At the same time, executive reporting remains weighted toward spending and operational activity rather than measures that directly connect engineering investments to business value.



Together, the findings point to a new challenge: AI is making engineering organizations more capable, but not smaller or cheaper. The CIOs who succeed in this era will be the ones who can identify which investments in talent and technology create measurable business value.

AI is creating capacity, not eliminating work

When software becomes cheaper to produce, companies want more of it. This dynamic is consistent with an economic concept known as Jevons paradox. When a resource becomes more efficient to use, total consumption can increase as new uses become economical.

As AI reduces the energy required to generate code, organizations can pursue more software projects and ship more features, consuming productivity gains through greater output rather than simply banking them as lower headcount.

This runs counter to the prevailing narrative that has emerged around AI over the past three years, which was roughly:



More Productivity → Fewer People → Lower Costs

In reality, the demand for more output (more code, more prototypes, more features, more products, etc.) is creating more work for software engineers, despite some reduction in time spent coding.



More Productivity → Flat Or Growing Headcount → More Output

How AI is impacting engineering output and headcount

66%

More output
same headcount


24%

More output
growing headcount

10%

Same output
fewer engineers

Change over the past year for most organizations



As AI lowers the energy required to create software, it expands the universe of software worth creating. Backlogs that seemed endless are now achievable. Ideas that previously could not justify an engineering team's time can now be prototyped. Internal tools become economical to build.

Rather than taking the productivity dividend primarily through lower headcount, organizations are reinvesting it in greater output.

“The biggest change I’m seeing with AI isn’t that companies need fewer engineers. It’s that the same teams can take on more. When you dramatically increase what an engineer can accomplish, the natural response isn’t necessarily to shrink the team. It’s to raise the ambition for what the team can build.”



John Burns, Former CIO

More output does not necessarily mean faster delivery

While 82% report increased product and feature volume, 67% report an increase in the time required to bring new features to market. And 44% report an increase in coding errors.

That does not mean AI is making individual engineers less productive. It suggests that making one part of a complex system faster does not automatically make the entire system faster.

As code becomes easier to produce, constraints can move elsewhere.

The result is that organizations are producing more software without eliminating the friction it takes to get that software to market.

“One of the things I’ve learned running large technology organizations is that removing one constraint usually exposes another. AI can dramatically increase the capacity of individual engineers, but the rest of the organization doesn’t automatically accelerate with them. As engineering capacity grows, the bottleneck may move to product, governance, testing or decision-making. That’s why more engineering output doesn’t necessarily translate directly into faster business outcomes.”




John Burns, Former CIO

Less coding does not mean less engineering

The way engineers work is also evolving.

As AI takes on more coding responsibilities, engineers are spending more time reviewing AI-generated outputs, gathering and refining requirements, designing systems, testing, and managing the risks surrounding increasingly abundant code. That redistribution of engineering work up the stack requires a more sophisticated (sometimes perceived as “more senior”) skillset.



Less coding does not mean less engineering (continued)

As implementation becomes cheaper, these higher-order activities become a larger part of the engineering value chain.

This changes more than individual jobs.

Mala Singh, Chief People Officer at Electronic Arts, describes the organizational challenge as “pixelating” work: breaking a role into its underlying tasks and capabilities, identifying which pieces AI can absorb, and reconstructing the role around the human capabilities that remain most valuable.

“As AI absorbs more of the execution work, the skills that differentiate people move higher up the stack. Critical thinking, judgment, synthesis, and adaptability become more important. The challenge for companies is understanding which capabilities will actually create value in this new environment, and making sure they have the right people in the right roles.”



Mala Singh, Chief People Officer, Electronic Arts

Do that across enough roles and the structure of the engineering organization itself begins to change. The Human + AI workforce cannot simply be today's engineering organization with AI tools added on top. It has to be redesigned around a new division of labor between people and machines.

The challenge is that the division of labor itself keeps moving. David Lau, Vice President of Workload Scaling at OpenAI, described model capabilities advancing quickly enough that ways of working can become outdated within months.

“Frontier model capabilities are advancing so quickly that a task a model struggled with last month may be a one-shot success today. Engineers and leaders need to continuously re-evaluate what models can do and reassess how their teams work, so that both humans and agents are deployed where their relative strengths create the most leverage.”



David Lau, Vice President
Workload Scaling, OpenAI




AI is stretching the engineering performance curve


AI was sometimes described as a great equalizer: give every engineer access to the same powerful tools, and the differences in individual productivity should narrow.

Technology leaders increasingly believe the opposite.

The value of a strong engineer is rising (Global)



The upper end of the curve is stretching (U.S.)



Globally, **75%** now say a strong engineer is worth at least **3x total compensation**. That is up from **45%** in 2019.

The upper end of the curve is stretching even faster. In the U.S., **29%** now say a strong engineer is worth at least 5x total compensation, compared with **18%** in 2025 and **11%** in 2019.

As AI takes on more execution, the capabilities that differentiate engineers increasingly include judgment, systems thinking, domain expertise, product understanding, problem framing, and the ability to direct and evaluate AI.

“The most impactful engineers are increasingly applying skills we once associated primarily with engineering leaders. They’re constantly asking how best to combine the strengths of AI agents and human collaborators, structuring work so that each multiplies the impact of the other, and defining success criteria and work products that let agents demonstrate—clearly and verifiably—that they’ve achieved the desired outcome.”




David Lau, Vice President Workload Scaling, OpenAI

The engineering organization is beginning to reflect that shift

Organizations are investing more in senior talent.

Senior IC roles are growing while entry-level roles face greater pressure. Upskilling is also increasing as organizations try to develop AI-era capabilities within their existing workforce.

As a result, the traditional engineering pyramid may be changing to more of a diamond shape.



That can create near-term leverage by concentrating more work among experienced engineers equipped with AI. But it raises an important longer-term question: If AI absorbs more of the work that historically developed junior engineers, where will the next generation of senior engineers come from?

AI is also raising the bar for IT Service Providers

The premium on talent quality extends beyond the full-time workforce.

Leaders are recognizing the quality gap between their full-time engineers and contractors: 57% of U.S. CIOs say they are very satisfied with FTE quality, compared with just 38% who say the same about contractor quality.

The talent quality gap becomes more consequential as AI increases the productive capacity of every engineer. In theory, AI should improve the economics of outsourced engineering. More productive contractor teams could produce more for the same spend, or allow clients to achieve the same output at lower cost.

So far, U.S. leaders predominantly expect the former. 51% expect contractors to deliver more output at roughly the same cost, mirroring what they're seeing out of their own organizations.

That creates a new test for the contractor value proposition. If AI-driven productivity does not substantially lower client costs, IT Service Providers will have to differentiate on the quality of their talent and the additional output that talent can produce.

AI readiness is an economic differentiator

The importance of talent quality also appears in how organizations experience AI returns.

Among organizations where technology employees are considered very prepared to leverage AI in their day-to-day work, 89% say at least half of their engineers are seeing meaningful ROI from AI tools. Among organizations whose employees are not considered very prepared, that falls to 65%.

That is a 24-point gap in reported AI ROI.

This finding reinforces a broader pattern throughout the research: Putting powerful AI tools in front of engineers is not enough. The capability of the people using those tools matters.

“The question for engineering leaders shouldn’t just be, ‘How much AI are our engineers using?’ It should be, ‘Are they using it well?’ The engineers who create the most value will be the ones who know how to direct AI, evaluate what it produces, and make the right decisions with it in the loop.”



Sagnik Nandy, CTO, Docusign

Compensation models are starting to feel the pressure

If AI increases the leverage of the best engineers, a natural question follows: Will compensation eventually reflect the widening difference in the value they create?


So far, most organizations are adapting within existing compensation structures. In the U.S., 78% say top-performing engineers receive higher compensation, but within established pay bands. Only 15% say their organizations pay top performers significantly more in ways that meaningfully reflect their business impact or ROI.

That creates a widening gap between how organizations value differentiated engineering talent and how they pay for it. Nearly three in ten U.S. leaders now believe a strong engineer can be worth at least 5x their total compensation, yet most organizations are still rewarding that additional leverage within compensation systems designed for a narrower performance curve.

As AI continues to stretch the engineering performance curve, compensation systems will eventually have to stretch with it. 82% of U.S. technology leaders expect compensation differences between top-performing and lower-performing engineers to increase as AI further amplifies the impact of their best engineers.

The challenge will be determining which engineers are actually creating that differentiated value and measuring it consistently enough to reward it.

The total costs of Human + AI engineering are changing faster than the ability to benchmark and measure the impact



Everyone can't be ahead of the market

The difficulty in understanding Human + AI ROI begins with a surprisingly basic question: What does good AI performance actually look like?

Nearly three-quarters of U.S. technology leaders believe their organizations are ahead of their peers with respect to AI. 13% say they are significantly ahead, and 61% say they are slightly ahead. Another 21% believe they are in line with the market. Just 5% believe they are behind.

The overwhelming perception of overperformance reveals the challenge leaders face in assessing their AI maturity. If leaders cannot reliably determine whether their engineering organization is ahead of the market, it becomes harder to know whether additional AI spend or workforce spend is producing a competitive advantage or merely keeping pace.

AI is changing what engineers do and what they cost

Traditional workforce planning treated software engineers as a human-capital expense: salary, benefits, equipment, and overhead. AI adds a new variable component: models, agents, tokens, coding tools, and compute.

Engineering leaders are already beginning to incorporate these costs into workforce decisions. In the U.S., 75% say they consider an engineer's total cost, including AI usage, when making workforce decisions. Another 74% say they actively optimize AI spending at the individual engineer level. Thus, the economics of an engineer can no longer be understood from compensation alone.

“AI is changing the total cost of an engineer. Salary is no longer the whole equation. You’re increasingly adding models, agents, and token consumption on top of it. Before you give every engineer effectively unlimited access to those resources, you need confidence they have the skills and judgment to use that leverage well.”



Sagnik Nandy, CTO, Docusign

Measuring cost is easier than measuring value

The same uncertainty appears in executive reporting. The emerging equation looks something like this:




Human cost + AI cost → engineering output → business value


Organizations are getting increasingly sophisticated at measuring the left side of that equation. The right side is harder.

How do you evaluate and measure the value of engineers?

* = business impact



What metrics do you report to your CEO/CFO



Technology leaders track a range of spending, adoption, and operational measures. But measures that directly connect engineering investment to business outcomes remain less prevalent.

That distinction becomes more important as the engineering cost base changes. Organizations are largely maintaining or growing headcount. They are increasing investment in senior talent and upskilling. Their strongest engineers are perceived as increasingly valuable. And AI adds another layer of variable cost. Today, this is driving up costs for engineering organizations, while also making them more capable.

But the question that CEOs and CFOs will be asking next is, “What did we get for that incremental investment?”

Business ROI will depend on how the added engineering capacity is deployed, what it produces, and how it ultimately benefits the business.

Building the Human + AI engineering organization

The research points toward a different management challenge than the one technology leaders faced during the first wave of generative AI. Three capabilities become particularly important.

1. Understand how work is changing

Organizations need to continually break engineering work into its underlying tasks and capabilities, understand which activities AI can absorb, and determine which human capabilities become more valuable as a result.

2. Identify and deploy differentiated talent

When every engineer has access to increasingly powerful AI, coding ability alone stops being a differentiator. Judgment does not.

Organizations need to understand which engineers can frame ambiguous problems, provide context, evaluate AI output, make architectural decisions, and connect technical work to business outcomes.

Those capabilities should inform hiring, contractor selection, internal mobility, project assignment, development, and compensation.

“More AI usage isn’t the goal. More code isn’t the goal. The goal is to build an organization where the combination of engineers and AI is creating more value for the business. Whether you can measure that well enough to know where to invest next will be the key to success.”



Sagnik Nandy, CTO, Docusign

3. Benchmark technology and talent against business value

Organizations need a clearer picture of both sides of the Human + AI equation.

- Are they spending more or less than comparable organizations on AI?
- Are their engineers actually more capable of using it?
- Are contractors delivering enough additional value to justify their cost?
- Which engineers create disproportionate leverage when paired with AI?
- And most importantly, does that combination of technology and talent create measurable business value?

AI adoption and engineering activity are useful measures of what an organization is putting into the system and what it is getting out. But neither tells leaders whether those investments are creating meaningful business value. The real challenge is understanding how technology and talent work together: Which engineers create differentiated leverage with AI, what that combined capability costs, and where deploying it produces returns that justify the investment.

Conclusion

The Human + AI era is raising the stakes on engineering talent quality.

AI is making software easier to produce, and CIOs are responding by producing more of it.

The productivity dividend is not showing up as smaller engineering organizations. AI is increasing engineering capacity and shifting human work toward higher-order activities.

In this environment, the strongest engineers become disproportionately more valuable. Talent quality becomes more consequential across both employees and contractors. AI adds variable technology costs to the workforce.



“We used to evaluate engineers largely by the code they wrote, how much they produced, how quickly they produced it, and how well it worked. As agents take over more of the execution, we’ll increasingly evaluate engineers by the quality of the questions they ask and how they navigate uncertainty: how quickly they can turn a broadly scoped, ambiguous engineering challenge into a sequence of increasingly precise questions, decisions, and experiments that illuminate an increasingly clear path toward the best solution.”



David Lau, Vice President Workload Scaling, OpenAI

In this environment, the strongest engineers become disproportionately more valuable. Talent quality becomes more consequential across both employees and contractors. AI adds variable technology costs to the workforce.

At the same time, leaders are trying to manage the new economics without clear external benchmarks. The organizations that lead the new era will be those that understand which technology and talent capabilities produce business value.

AI is making engineering organizations more capable. The advantage will come from knowing what that capability is worth.