AI Workforce Transformations:

Financial Services

karat

Executive Summary

AI is transforming how engineers work. Yet, many financial services firms are still approaching their engineering workforce strategy with a pre-AI playbook, increasing the risk of falling further behind their tech-industry peers.

For CIOs and CTOs, the question is whether existing talent strategies, hiring practices, and quality standards are sufficient to safely and effectively scale AI across mission-critical systems.

As AI accelerates the speed and complexity of engineering work, CIOs who modernize how they identify and enable the AI-ready engineers will be best positioned to compete in the human + AI era of software engineering.

Key findings

- AI disproportionately amplifies the output (and risk) across teams and individual engineers. The strongest engineers generate the majority of AI-driven productivity gains, yet elite engineering talent remains concentrated in tech companies.
- AI is already being deployed at scale across most engineering organizations, with paired programming as the most common use case.
- Despite cost pressure, CIOs and CTOs expect headcounts to remain flat or increase over the next three years, raising the stakes on hiring quality.
- Financial services organizations trail tech peers in both engineering quality and the ability to assess AI-native skills.
- Human + AI collaboration is emerging as a critical differentiator in measuring engineering performance.

How AI is transforming software development

As expected, most organizations are using AI primarily for paired coding. Testing/QA, code review, data science, vulnerability detection, and agentic AI round out the top use cases, with leaders identifying agentic AI as delivering the highest ROI.

Tech companies are currently realizing slightly higher productivity gains from AI than financial services firms (36% vs. 32%). However, several indicators suggest this gap is likely to widen, particularly as AI moves from isolated productivity tooling into system-level automation and decision support.

Importantly for financial institutions, leaders also anticipate a significant increase in legacy modernization and refactoring over the next 3—5 years. This aligns with the reality that many core banking, payments, and risk systems were not designed for AI-accelerated development. This work also requires a more nuanced and human-led AI touch. While moving from monolithic systems to micro-architectures and agent-driven workflows makes generating v1 is easier, durability, auditability, and safe change remains the real bottlenecks. The longer a piece of code needs to last, the more human oversight, architectural thinking, and explainability it needs.

Engineering quality is more important than ever

The gap between strong and weak engineers is growing

73% Of US leaders agree strong engineers are worth at least 3x their total comp

59% Say weak engineers deliver net-zero or negative value

Engineering leaders disproportionately attribute the gains from AI to their strongest engineers. Almost three-quarters of leaders in the U.S. now say that strong engineers are worth at least 3x their total compensation, up from 45% just five years ago.

In financial services, where defects, outages, and security failures carry outsized consequences, this widening gap has implications beyond productivity. AI magnifies both excellence and error. Without strong engineering judgment, AI can accelerate the wrong outcomes just as efficiently as the right ones.

Pressure to cut costs and low engineering quality are throttling AI gains

While AI promises faster development cycles and expanded feature delivery, many financial services organizations risk sacrificing long-term value in favor of short-term cost optimization.

We see tech companies more aggressively investing and implementing AI tools at scale. Meanwhile finserv leaders are being asked to cut cotsts due to gains before they've had a chance to realize those gains. This correlates with more tech leaders feeling their teams are "very prepared to leverage AI day to day."

For CIOs, this tension is nothing new. CFOs push for efficiency, while technology leaders are accountable for resilience, modernization, and long-term competitiveness. The data suggests that under-investing in talent readiness is already limiting AI ROI.

70% of finserv leaders reported feeling pressure to cut costs due to AI gains compared to **59%** of tech leaders.

% who are implementing AI tools at scale:

Tech: 61%; **Finserv:** 43%

% who say employees are "very prepared" to leverage AI day-to-day:

Tech: 60%; **Finserv** 39%

The AI workforce transformation gap

Contractors remain a critical component of financial services engineering capacity, with 87% of finserv leaders expecting contractor headcount to stay flat or increase in the coming years.

However, Karat's data shows that engineers from IT Service Providers (ITSPs) often underperform full time engineers at big banks and financial institutions. Without consistent benchmarking and quality standards, organizations risk embedding lower-quality engineering talent into core systems, precisely as AI increases the leverage of individual contributors.

For CIOs overseeing complex vendor ecosystems, this represents a growing blind spot: AI adoption without talent quality controls increases operational and technology risk.

Distribution of engineering talent quality by industry

The skillset required to be a "strong" engineer is growing

While strong engineers are producing more value than ever, the skills needed to drive that ROI are expanding.

While foundational skills like problem-solving, communication, and product sense remain constant, leaders are also seeking out new AI-native abilities, while skills like writing code without AI assistance become less relevant. When asked what skills leaders wanted to add to their organizations, the top-5 responses all related to AI:

1. **Familiarity with agentic AI**
2. **Using AI for coding**
3. **Integrating 3rd-party AI APIs**
4. **Prompt engineering**
5. **Evaluating and mitigating AI-related risks**

Talent measurement isn't keeping pace with AI

Ten years ago, a strong engineer could quickly apply the right data structure or algorithm to produce code that led to a working solution to a problem. Today, a strong engineer might still do that. But if they aren't using the latest tools, they will quickly fall behind a similarly skilled engineer who can leverage AI.

The problem is, most organizations aren't yet assessing AI skills. 62% of leaders don't allow candidates to use AI in interviews. Just 30% of organizations ranked "updating technical assessments to assess for AI" as a top priority, and even fewer, 25% are prioritizing training interviewers. To make matters worse, leaders estimate that over half of candidates are using AI despite being instructed not to.

The impact is twofold:

1. Leaders aren't getting signals on the AI skills they're looking for.
2. The signal from interviews that aren't AI-enabled is useless because most candidates are using AI.

This is a troubling blind spot because most organizations lack the expertise or resources to invest in their talent measurement systems. This problem is especially acute for financial services companies, where the pressure to cut costs has led leaders to focus more on efficiency than on improving engineering quality, which is often perceived as being harder to measure or control.

Engineering workforce priorities (ranked as top-2 priority)	Finserv Orgs.	Tech Orgs.
Improving the efficiency of the hiring process	47%	32%
Building an AI-ready Workforce	42%	41%
Raising my hiring bar	41%	43%
Updating Technical assessments to assess for AI skills	25%	33%
Replacing human workflows with AI	24%	25%
Training interviewers to assess AI skills	22%	27%

Next-generation interviews are essential for assessing AI skills

Pre-LLMs, it was possible to extrapolate a software engineer's ability to gather the appropriate information, consider edge cases, and optimize code based on having them present a working solution to a simple problem.

But in the human + AI era, producing a working solution no longer generates the same hiring signal. Interviews must now isolate and independently measure a candidate's underlying skills while also assessing AI readiness.

As a result, the engineering leaders who aren't investing in next-generation assessments are getting less predictive hiring signals, and confidence in hiring is eroding.

Increased difficulty assessing skills

% US leaders who say it is difficult to accurately assess skills needed to capitalize on new developments in AI

Decreased hiring confidence

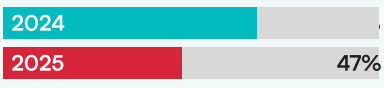

% US leaders who are very confident that qualified candidates are the ones getting job offers

AI to blame

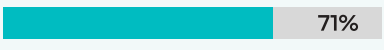

% leaders who say AI is making it harder to assess technical skills

It's clear that leveraging AI tools is now a huge part of the way most developers get their work done. We expect that trend to continue.

We also know that the best interviews get as close as possible to mimicking real-world work. A decade ago, that meant moving technical interviews from a shared Google Doc to a fully-tooled IDE. Today, it means allowing candidates to leverage LLMs in their interview experience while interacting with live engineers throughout the interview.

Organizations using this combination of human-led, AI-enabled interviews consistently anticipate stronger AI outcomes than those using human-only or non-human assessments.

Human + AI: 33% of organizations in the U.S. use live interviews and allow candidates to use AI tools in a way that mirrors day-to-day work

Over the next 3 years...

- **63%** of companies that use human + AI interviews expect coding errors to decrease
- 45% of companies that use human-only interviews
- 21% of companies who use non-human assessments

Human-only: 45% of organizations in the U.S. measure technical skills using live interviews where AI use is prohibited

- **49%** of companies that use human + AI interviews expect the time it takes to bring new products/ features to market to decrease
- 31% of companies that use human-only interviews
- 20% of companies who use non-human assessments

Non-human: 22% of organizations in the U.S. use a combination of online code tests and take-home projects that do not include human interviewers

- **76%** of companies that use human + AI interviews expect the number of products and features they release to increase
- 57% of companies that use human-only interviews
- 48% of companies who use non-human assessments

The challenge is that it takes a lot of time and effort to build & maintain effective AI-enabled interviews. You can't just drop an LLM into an IDE. Adding AI to interviews requires investments. Introducing a multi-file code base, adding LLMs to your IDE, rapidly iterating and developing new content as LLMs evolve, and constantly training interviewers are now table-stakes for technical hiring.

Adding AI to interviews requires a thoughtful approach, an investment in people and systems, and ongoing effort to improve as the technology and best practices evolve.

Key questions that CIOs should be asking themselves include:

- How can we assess engineers in a realistic environment?
- How do we offer a consistent AI-enabled environment for evaluations?
- How will we evolve rubrics and assessments as models and practices evolve?

The most successful organizations are addressing these questions head-on. Tech companies are already investing more time in talent acquisition by assessing more candidates per hire. Tech companies interview an average of 25.5 candidates per hire, with 16.8 making it to the onsite loop compared to 18.2 and 9.4 for finserv companies, respectively.

Compounding the challenge for finserv organizations, these same tech leaders are expecting to double down on their time spent on talent acquisition, with 43% of tech leaders saying they expect the time they spend on talent acquisition to increase over the next 3 years, compared to 24% of finserv leaders.

This is a very real acknowledgement of the value of engineering talent, which remains an organizaitons greatest asset in the human + AI era.

Building the Human + AI Workforce

AI is not reducing the importance of engineering talent; it is dramatically increasing it.

For CIOs and CTOs, the central challenge is no longer whether AI can drive productivity, but whether their organizations can safely and reliably scale that productivity across regulated, mission-critical systems. As AI amplifies individual impact, gaps in engineering quality, hiring signals, and workforce readiness become systemic risks that can reduce the overall value AI creates for the full organization.

The data is clear: organizations that combine human judgment with AI-native assessment practices are better positioned to reduce errors, accelerate delivery, and sustain innovation without sacrificing control. The next wave of software engineering will not be defined by fewer engineers or lower costs. It will be defined by higher-quality engineers working with increasingly powerful tools, under governance models that evolve as fast as the technology itself.

Whether it’s developing software or making the right talent decisions, Human + AI isn't a static configuration. It's a philosophy about how to navigate technological change. The future isn't human OR AI. It's the two working together.

Closing the AI talent gap in financial services requires a human + AI workforce strategy: one that raises the bar and continuously adapts to the realities of the AI era.