

2026

# State of Tech Talent Report

Not a jobs crisis,  
but a skills crisis with  
an upskilling answer

May 2026


Marco Gerosa, Ph.D., *Northern Arizona University*

Adrienn Lawson, *The Linux Foundation*

Anna Hermansen, *The Linux Foundation*

Foreword by Clyde Seepersad,

*Linux Foundation Education*



# 2026 State of Tech Talent Report

**97% plan to use AI,** and 55% expect significant value in software development.



**AI continues to expand technical hiring in IT,** with aggregated net hiring effects of +26% in 2025 and +31% in 2026.

**AI is NOT eating all IT jobs,** as only the largest organizations are reporting a negative net hiring effect (-4%).



**Understaffing remains widespread in AI (47%),** cybersecurity (40%), cost optimization (36%), and platform engineering (34%), though the trend is improving.



**Capability gaps in AI security** and AI operations each affect 57% of organizations.



**Security concerns (48%) and budget constraints (47%)** are the leading barriers to getting value from new technologies.



Upskilling existing staff is the leading strategy for **closing AI talent gaps and preserving institutional knowledge,** rated important by 94% of organizations.



**Upskilling existing staff (57%)** is the primary response to talent gaps, ahead of external hiring (49%).



Organizations are **3.5 times more likely to upskill than to hire** across strategic technological domains.



**Upskilling is favored over hiring** for business context (7.9x), retention (7.7x), team cohesion (7.3x), and lower total cost (5x).



**Technical training (93%) ranks above compensation (91%)** as a retention strategy, suggesting that technical talent values opportunities to grow.



**76% of hiring managers consider certifications important** for evaluating a candidate's technical skills.



# Contents

Foreword.....	4
Executive summary.....	6
Introduction.....	7
AI creates value, but is not eating all IT jobs.....	8
Gains are expected across core activities.....	8
AI and the future of IT jobs.....	9
AI reshapes technical roles.....	11
The skills gap is in the full stack, not just AI.....	13
Organizations are understaffed in AI and other strategic domains.....	13
Most organizations lag behind in PARK stack implementation.....	15
Agentic AI raises serious security risks, and most organizations are not ready.....	16
Security challenges are rising in relevance.....	16
Security concerns are the leading barrier to getting value from AI.....	17
The value of institutional knowledge.....	19
Upskilling is the preferred strategy to address the lack of a skilled workforce.....	19
Upskilling is preferred across the technology stack.....	21
Upskilling vs. hiring: benefits and challenges.....	22
Technical training is among the top strategies to retain technical talent.....	23
The factors organizations value most in technical assessments favor internal candidates.....	24
Conclusion.....	25
Methodology and demographics.....	27
About the survey.....	27
DataWorld access.....	28
Respondent demographics.....	28
Appendix.....	30
About the authors.....	37
Acknowledgments.....	38

# Foreword

I'm pleased to share the 2026 State of Tech Talent Report, our annual look at how organizations are building, developing, and deploying technical talent. As artificial intelligence continues to reshape both roles and expectations, the top-of-mind question is: Where are we headed, and what talent will we need to get there?

If there is one conclusion that stands out this year, it's this: the narrative around AI and jobs continues to miss the point.

Across our data, AI does not appear to be reducing demand for technical talent. It is raising expectations for what that talent must be able to do. Organizations are not simply pulling back on hiring, they are redefining roles in real time and expanding the scope of responsibilities across software development, infrastructure, data, and security. External research reinforces this shift. While AI adoption is now widespread, most organizations are still far from realizing its full value at scale.<sup>1</sup> The constraint is not access to technology, it is the ability to operationalize it effectively.<sup>2</sup>

That gap is most visible in how organizations are approaching risk. Unlike traditional software, AI systems introduce variability, opacity, and new failure modes that are difficult to predict and control. As a result, security, governance, and operational discipline have moved from secondary considerations to primary barriers.<sup>3</sup> Many organizations are discovering that deploying AI is not simply a tooling decision, it is an infrastructure, security, and skills challenge that cuts across the entire technology stack.

This is where the talent conversation becomes more urgent and more nuanced.

For years, the industry has framed the “talent shortage” as a question of supply. This year’s data suggest a different reality. The issue is not simply a lack of people. It is a lack of capability—specifically, the ability to apply skills across systems, workflows, and real-world environments. Organizations are struggling not

---

1 McKinsey & Company, *The State of AI 2025: Agents, Innovation, and Value Creation*, 2025.

<https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>

2 Boston Consulting Group, *Where's the Value in AI? Only 5% of Companies Are Getting It at Scale*, 2025.

<https://www.bcg.com/publications/2025/ai-value-creation>

3 Deloitte, *State of Generative AI in the Enterprise*, updated 2025 findings. <https://www2.deloitte.com>

because they cannot hire, but because they cannot easily find or develop professionals who can operate AI-driven systems reliably and securely in production.<sup>4</sup>

In response, leading organizations are changing how they think about talent sourcing. Rather than relying solely on external hiring, they are investing in upskilling their existing teams, often because it is faster, more effective, and more aligned to their environments.<sup>5</sup> Besides, upskilled talent doesn't take wing and fly, they grow roots and thrive. This reflects a broader shift we see across the industry: workforce transformation is becoming inseparable from technology transformation.

At the same time, the way organizations evaluate talent is evolving. Traditional signals, such as degrees, titles, and years of experience, are proving insufficient in a market where roles are changing faster than those signals can keep up. Employers are placing greater weight on demonstrated, job-relevant skills and on credible ways to validate them. Industry-recognized certifications and hands-on experience are increasingly used to establish trust in what professionals can actually do, not just what they have studied.<sup>6</sup>

These trends point to a fundamental shift in how technical talent is defined and developed. AI is not creating a shortage of opportunity. It is redefining where opportunity lives and raising the bar for those who want to seize it.

The organizations that succeed in this environment will not be those that simply adopt AI tools. They will be the ones who invest in building the skills, structures, and validation mechanisms needed to use those tools effectively. In the end, the advantage will not come from technology alone, but from the people who know how to apply it.

### **Clyde Seepersad**

Senior Vice President and General Manager, Education, Linux Foundation

---

4 Bain & Company, Widening Talent Gap Threatens Executives' AI Ambitions, 2025. <https://www.bain.com>

5 Boston Consulting Group, AI Transformation Is a Workforce Transformation, 2026. <https://www.bcg.com>

6 LinkedIn, Skills on the Rise / Skills-Based Hiring Trends, 2025. <https://www.linkedin.com>

# Executive summary

Over the past several years, the Linux Foundation has surveyed hiring and training stakeholders to capture the state of the technical talent market amid technological shifts and economic changes. This year's study is based on an online survey fielded in February 2026, which collected responses from 400 participants worldwide and examined the impact of AI, especially generative AI, on the IT talent market.

The results suggest that the primary obstacle to realizing value from AI is not the technology itself, but a growing security and operational readiness crisis. Organizations are accelerating AI adoption (97% plan to implement AI), but are failing to build the foundational capabilities necessary to deploy it safely. Security and privacy concerns have risen sharply to become the top barrier to adopting new technologies, increasing from 17% in 2024 to 48% in 2026.

This risk is symptomatic of a broader full-stack readiness problem. The report shows that organizations are understaffed across strategic domains, and these staffing shortages are associated with capability gaps. A majority of organizations report capability gaps in AI security and risk management (57%) and in AI operations and monitoring (57%). Operational maturity is further constrained by the low deployment of fundamental AI infrastructure, indicating that most organizations are still building the capabilities required for production AI.

Despite negative media reports about layoffs in the tech industry due to AI, our findings show that AI is acting as a net driver of job creation in IT, with a net hiring effect of +31% expected for 2026. The data suggest that recent layoffs are concentrated at the largest enterprises (20,000+ employees), the only segment reporting a negative net hiring effect (-4%). Smaller organizations and end-user companies may be absorbing displaced technical talent, as both report strong positive net hiring effects. Demand is particularly high for AI-specific roles, which show a net hiring effect of +60%.

To bridge these full-stack and operational gaps, organizations prefer to look to their internal staff. Upskilling and cross-skilling existing staff is the top strategy (57%), favored over external hiring (49%). This approach offers major advantages in preserving institutional knowledge and is strongly preferred for understanding business context (7.9x) and staff retention (7.7x). Hiring externally, by contrast, is slower and riskier: new hires take 53% longer to reach productivity, and 28% resign within six months. Overall, the findings suggest that technical professionals value learning and development at least as highly as compensation when deciding whether to stay. The path to sustainable AI scaling requires prioritizing internal development and continuous learning to build secure, operational full-stack competencies.



# Introduction

There is little doubt that AI is changing the IT job market. Adoption is widespread and organizations across industries and regions are integrating AI into core business functions. This widespread adoption also raises important questions: Where do organizations expect AI to deliver value? Is that value coming at the cost of technical jobs, or is it generating new ones? What skills does operationalizing AI actually require, and how are organizations building them?

**Organizations need to know which technical domains to prioritize, whether to hire or train staff, and how to close skills gaps that go beyond AI.** At the same time, technical talent needs to understand how to position themselves for emerging career opportunities.

This report addresses those questions, drawing on survey responses from 400 global participants responsible for hiring, training, and managing technical talent. It examines where organizations expect AI to create value, how AI is reshaping the technical workforce, which skills are needed to support AI in production, and how organizations are responding to the talent gaps.

Our analysis highlights the value of skill development programs, including upskilling (deepening existing capabilities) and cross-skilling (expanding expertise across domains). For consistency, we use the term “upskilling” throughout this report to refer to both of these training approaches.

# AI creates value and is not eating all IT jobs

## Gains are expected across core activities

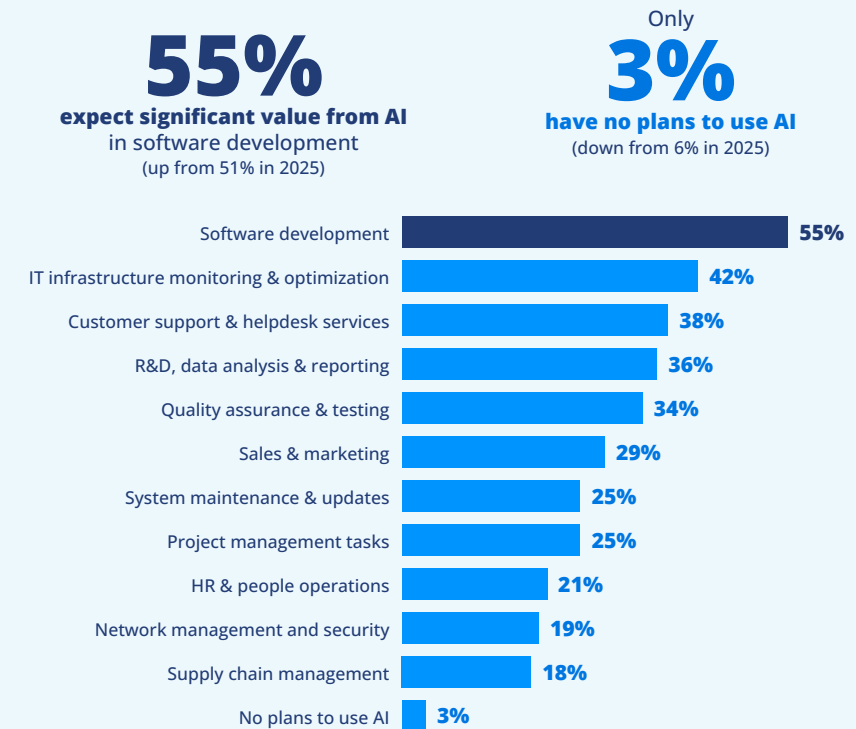
AI has become hard to avoid. In our survey, only 3% of organizations have no AI plans, down from 6% in the 2025 survey<sup>7</sup>. This widespread adoption is corroborated by other surveys, including one conducted by McKinsey in 2025,<sup>8</sup> which found that nearly 90% of organizations regularly use AI in at least one business function. Our survey also reveals that organizations expect to obtain value from AI across a broad range of core activities, as shown in Figure 1.

Software development (55%) ranks first among the surveyed AI value drivers, consistent with last year's findings. Organizations also expect substantial value in related activities such as quality assurance and testing (34%) and system maintenance and updates (25%). These results are consistent with prior research on productivity gains from AI-assisted development<sup>9,10</sup> and recent industry investments, including the Linux Foundation's launch of the Agentic AI Foundation to support the emerging ecosystem of AI coding agents and agentic infrastructure.<sup>11</sup>

**FIGURE 1: EXPECTED AI GAINS ACROSS CORE ACTIVITIES**

Where does your organization expect AI to deliver significant value? (select all that apply)

2026 Tech Talent, Q23, Sample Size = 400, Total Mentions = 1,378, DKNS excluded (3%)



7 Gerosa and Lawson, "2025 state of tech talent: Truth vs. vibe—the not so disruptive workforce impact of AI," *The Linux Foundation*, Jun. 2025. [Online]. Available: <https://www.linuxfoundation.org/research/open-source-jobs-report-2025>

8 Singla et al., "The state of AI in 2025: Agents, innovation, and transformation," *McKinsey & Company*, Nov. 2025. [Online]. Available: <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>

9 Cui et al., "The effects of generative AI on high-skilled work: Evidence from three field experiments with software developers," *Management Science*, 2026. <https://pubsonline.informs.org/doi/10.1287/mnsc.2025.00535>

10 Peng et al., "The impact of AI on developer productivity: Evidence from GitHub Copilot," arXiv:2302.06590, 2023, <https://arxiv.org/abs/2302.06590>

11 Linux Foundation, "Linux Foundation announces the formation of the Agentic AI Foundation," *Press Release*, Dec. 9, 2025. <https://www.linuxfoundation.org/press/linux-foundation-announces-the-formation-of-the-agentic-ai-foundation>



## AI and the future of IT jobs

AI's productivity gains have long raised fears of AI-driven job displacement. These fears align with recent press reports attributing large-scale layoffs to AI.<sup>12</sup> According to outplacement firm Challenger, Gray & Christmas, AI-related reasons were cited in 54,836 job cuts in 2025, a year in which total U.S. job cuts reached their highest level since the COVID pandemic.<sup>13</sup>

Our survey of the technical talent market shows a different picture. Last year's survey suggested that, within the IT sector, more organizations were increasing technical headcount due to AI than decreasing it, with a net hiring effect of +18% in 2024 and expectations of +21% in 2025 and +23% in 2026.

The 2026 survey (Figure 2) suggests that reported outcomes exceeded those expectations. Organizations retrospectively report a net hiring effect of +26% for 2025, five percentage points above expectations, and +31% for 2026, eight percentage points above the original expectation. Growth is expected to slow in 2027, with the net hiring effect declining to +22%, a drop of 9 percentage points from 2026. Even so, the net effect remains positive across all three years.

Almost half of respondents report growing their technical workforce in response to AI-related demand, suggesting that AI is not (yet) displacing technical talent at the scale many anticipated. At least among surveyed organizations, AI currently appears to be acting more as a driver of demand than as a reducer of headcount. Other analyses are also optimistic. The World Economic Forum's *Future of Jobs Report 2025* projects a net global increase of 78 million jobs by 2030, with 170 million new roles created against 92 million displaced,<sup>14</sup> and an analysis by the Information Technology and Innovation Foundation found that AI-related activity generated approximately 119,900 direct jobs in the United States in 2024 alone, substantially outpacing AI-driven losses.<sup>15</sup>

---

12 Bhaimiya, "AI Was Behind Over 50,000 Layoffs in 2025 — Here Are the Top Firms to Cite It for Job Cuts," *CNBC*, Dec. 21, 2025. <https://www.cnbcm.com/2025/12/21/ai-job-cuts-amazon-microsoft-and-more-cite-ai-for-2025-layoffs.html>

13 Challenger, Gray & Christmas, "2025 Year-End Challenger Report: Highest Q4 Layoffs Since 2008; Lowest YTD Hiring Since 2010," Jan. 8, 2026. <https://www.challengergray.com/blog/2025-year-end-challenger-report-highest-q4-layoffs-since-2008-lowest-ytd-hiring-since-2010/>

14 World Economic Forum, "The Future of Jobs Report 2025," Jan. 2025. <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>

15 Ostertag, "AI's Job Impact: Gains Outpace Losses," Information Technology and Innovation Foundation (ITIF), Washington, DC, USA, Dec. 18, 2025. <https://itif.org/publications/2025/12/18/ais-job-impact-gains-outpace-losses/>

**FIGURE 2**

### A positive net growth in IT jobs from 2025 through 2027

2026 Tech Talent, Q21, Sample Size = 400, not applicable and DKNS excluded (6% in 2025, 13% in 2026, 30% in 2027). Net Hiring Effect = % increase - % decrease




The net hiring effect by organization size, region, and industry type (Figure 3) reveals that organizations with 20,000 or more employees are the only segment reporting a negative net hiring effect (-4%). All smaller segments report positive effects ranging from +28% to +42%, suggesting that the headline layoffs may be concentrated in large enterprises.

Some of the displaced talent may have been absorbed in smaller organizations or end-user organizations (those that leverage IT to offer products and services in other industries), which report a higher net hiring effect (+31%) than IT industry organizations (+22%). These organizations may need to expand their IT workforce to implement AI in their business models and internal processes. Regionally, we also observe some differences, with Asia-Pacific leading at +34%, followed by Europe at +23% and the US or Canada at +18%.

**FIGURE 3**

### Net hiring effect in 2025 across organization sizes, regions, and types

2026 Tech Talent, Q21 by Q8, Q9, Q13, Sample Size = 400, Net Hiring Effect = % increase - % decrease, detailed data available in Appendix A1



Our findings raise questions about the misalignment between our data and press reports attributing layoffs to AI. Macroeconomic conditions, post-pandemic workforce corrections, and geopolitical pressures are all plausible contributors to recent job cuts, factors that tend to attract less scrutiny when AI provides a more compelling explanatory frame. Nevertheless, our findings require careful interpretation since the survey captures self-reported perceptions within the IT sector and does not isolate AI's contribution from broader economic conditions, industry growth cycles, or organizational investment decisions that may have coincided with AI adoption.

It is also possible that aggregate headcount is declining in some segments while growing in others, and that the organizations represented here skew

toward those actively investing in AI. The size breakdown in Figure 3 reinforces this possibility as large enterprises with 20,000 or more employees are the only segment reporting a negative net hiring effect, while all smaller segments report strongly positive effects.

One interpretation is that technical talent displaced from large organizations is being absorbed by smaller firms, producing a redistribution of where technical work is performed. Finally, whether the findings presented here reflect a durable effect or a transitional phase that may correct as automation matures remains an open question.

## AI reshapes technical roles


We also observe a positive net hiring effect at the role level (Figure 4). AI-specific roles show the largest net hiring effect, at +60%, up 6 percentage points from the previous year, reflecting increasing demand for specialized AI expertise. Software development (+28%), technical management (+22%), IT operations (+17%), QA and testing (+16%), and entry-level technical roles (+8%) all increased relative to the previous year's survey, suggesting that AI is generating broad demand across the technical workforce. Even entry-level roles, often considered most vulnerable to automation, show a positive net hiring effect of +8%, a two-percentage-point increase over the previous year.

The high net hiring effect in software development is particularly notable as this is also the area where organizations report the highest expected value from AI (Figure 1). The U.S. Bureau of Labor Statistics projects that employment of software developers will grow 17.9% between 2023 and 2033, much faster than the 4% average across all occupations.<sup>16</sup>

**FIGURE 4**

### The net hiring effect was positive across technical positions

2026 Tech Talent, Q22, Sample Size = 400, not applicable and DKNS excluded (11% to 15%), 2025 Tech Talent, Q23, DKNS excluded (11% to 15%), Sample Size = 556. Net Hiring Effect = % increase – % decrease, detailed data available in Appendix A2



16 Machovec et al., "Incorporating AI impacts in BLS employment projections: occupational case studies," *Monthly Labor Review*, U.S. Bureau of Labor Statistics, Feb. 2025. <https://doi.org/10.21916/mlr.2025.1>

This growth coexists with the reshaping of roles within the profession. A Federal Reserve paper finds that coder employment continued to grow but decelerated sharply after the introduction of generative AI.<sup>17</sup> Developers themselves report that they expect their roles to change significantly, with greater emphasis on architecture, integration, and AI-enabled decision-making.<sup>18</sup>

This evidence suggests a labor market in transition. Organizations are hiring due to AI and demand is shifting toward more experienced and AI-capable profiles.

---

17 Crane & Soto, “AI and coder employment: Compiling the evidence,” Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System, Mar. 2026. <https://www.federalreserve.gov/econres/feds/ai-and-coder-employment-compiling-the-evidence.htm>

18 BairesDev, “Dev Barometer 2025: A New Generation of AI-Native Developers — The Impact of AI on Software Developers and What to Expect in 2026,” White Paper, 2025. [https://www.bairesdev.com/files/Dev\\_Barometer\\_2025\\_White\\_Paper.pdf/](https://www.bairesdev.com/files/Dev_Barometer_2025_White_Paper.pdf/)

# The skills gap is in the full stack, not just AI

## Organizations are understaffed in AI and other strategic domains

Even as hiring expands, organizations report significant understaffing across all surveyed technical domains (Figure 5). The shortage is not concentrated in AI specialization but reflects a broad deficit in the infrastructure and operational support upon which AI deployment depends.

Deploying AI agents or models in production requires robust platform engineering, container orchestration, cost governance, and security hardening. In other words, the skills gap is not just about knowing how to use AI, it is also about having the engineering knowledge required to deploy it.


AI and ML engineering leads at 47% understaffing, which is high but represents a 12-point improvement from 59% in 2025. Cybersecurity and compliance (40%), FinOps and cost optimization (36%), platform engineering (34%), and cloud computing (29%) all show similar patterns, with high levels of understaffing but downward trends, suggesting that organizations are making progress, however gradual. Across all five domains, perceived understaffing fell by an average of 14% between 2025 and 2026. This may suggest that organizations recognize the importance of these domains and have begun investing in upskilling their staff.

Staffing gaps inevitably lead to capability gaps in key AI competencies (Figure 6). 57% of organizations report capability gaps in AI security and risk management, and the same percentage report gaps in AI operations and monitoring, followed closely by cost optimization for AI workloads at 54%. AI infrastructure expertise (45%), change management and communication (43%), data management (40%), infrastructure and platform expertise (36%), and business domain expertise (35%) also show gaps for implementing AI.

**FIGURE 5: UNDERSTAFFING REMAINS HIGH ACROSS TECHNICAL DOMAINS, BUT HAS DECREASED SINCE 2025**

**% of organizations rating themselves as understaffed (significantly or somewhat)**

2026 Tech Talent, Q19, Sample Size = 400, DKNS excluded (3% to 14%), 2025 Tech Talent, Q19, Sample Size = 556, DKNS excluded (3% to 15%), detailed data available in Appendix A3



**14%** average decrease from 2025

These findings challenge a common assumption about AI readiness. Organizations often treat AI readiness as a matter of training staff on prompt engineering or new tools. However, the data suggest that there is also a need to build technical fluency and operational readiness across the full stack that supports AI in production.

FROM THE EXPERTS



Mumshad Mannambeth  
Founder & CEO, KodeKloud



KodeKloud

The biggest skill gap today is not access to information, but the ability to turn knowledge into capability. As AI accelerates change in tools, platforms, and roles, the organizations that thrive are the ones that make learning continuous. Upskilling must go beyond theory and focus on applied learning that prepares people for real-world work in an AI-powered environment.


At KodeKloud, we support this journey by helping individuals and teams gain that advantage through immersive, hands-on learning rooted in real engineering challenges. The result is greater confidence, stronger capability, and a faster path to meaningful impact on the job.

In other words, the skills gap is not an AI problem narrowly defined, but a full-stack readiness problem. This gap can constrain the value organizations expect AI to deliver, regardless of how much they invest in AI tooling alone.

FIGURE 6: ORGANIZATIONS REPORT CAPABILITY GAPS ACROSS KEY AI COMPETENCIES

% of organizations rating themselves as having a significant or minor capability gap

2026 Tech Talent, Q25, Sample Size = 400, excluding Not Applicable (7% to 10%), Gap = % reporting significant or minor capability gap, detailed data available in Appendix A4




# Most organizations lag behind in PARK stack implementation

Our results also show that the infrastructure required to deploy AI in production is adopted unevenly across the surveyed organizations (Figure 7). The PARK stack, a layered architecture for production AI first introduced by Ben Lorica,<sup>19</sup> is in early-to-mid stages of adoption, with deployment rates varying considerably across its four layers. Much like the LAMP stack bundled complementary technologies for building, hosting, and maintaining web applications, PARK provides a common shorthand for the integrated set of capabilities organizations must assemble to move AI from experimentation into production.

**FIGURE 7**

## Deployment rates for PARK infrastructure technologies, at least partially deployed

2026 Tech Talent, Q20, Sample Size = 400, DKNS excluded (3% to 15%), 'At least partially deployed' = fully + partially deployed, detailed data available in Appendix A5



The PARK stack comprises PyTorch and equivalent ML frameworks (**P**) for model development and refinement (which 47% of organizations have at least partially deployed), AI foundation models (**A**) for providing base intelligence (50%), Ray and distributed compute (**R**) for scaling applications across systems (29%), and Kubernetes and containers (**K**) for resource provisioning and orchestration (66%).

Distributed compute engineering, such as Ray, trails the PARK stack at 29%, suggesting that organizations are not yet scaling AI workloads or handling the associated systems-level challenges. AI foundation models and PyTorch-based ML frameworks sit in the middle, at 50% and 47%, respectively. Kubernetes and containers lead the stack at 66%, consistent with their maturity in production environments.<sup>20</sup>

Nevertheless, even at the high end, one-third of organizations have not yet adopted Kubernetes and containers for AI workloads. For inference use cases, the primary workload for most organizations, Kubernetes can support autoscaling, cost optimization, multi-provider fallback, and latency management, making it a natural fit as AI moves into production. Recognizing this, the Cloud Native Computing Foundation (CNCF) launched the Certified Kubernetes AI Conformance Program in November 2025,<sup>21</sup> a community-led effort to define and validate standards for running AI workloads on Kubernetes.

The low deployment of the PARK stack makes it clear that for most organizations, maturity is still being built. Expanding headcount and adopting AI tools are necessary steps, but they are not enough without the full-stack capabilities required to deploy and scale production AI.

19 Ben Lorica, The PARK Stack Is Becoming the Standard for Production AI, <https://gradientflow.com/park-stack/>

20 Lawson and Sica, "CNCF Annual Cloud Native Survey: The Infrastructure of AI's Future," Cloud Native Computing Foundation (CNCF), The Linux Foundation, Jan. 2026. [https://www.cncf.io/wp-content/uploads/2026/01/CNCF\\_Annual\\_Survey\\_Report\\_final.pdf](https://www.cncf.io/wp-content/uploads/2026/01/CNCF_Annual_Survey_Report_final.pdf)

21 Cloud Native Computing Foundation, "CNCF Launches Certified Kubernetes AI Conformance Program to Standardize AI Workloads on Kubernetes," CNCF, Nov. 11, 2025. <https://www.cncf.io/announcements/2025/11/11/cncf-launches-certified-kubernetes-ai-conformance-program-to-standardize-ai-workloads-on-kubernetes/>

# Agentic AI raises serious security risks, and most organizations are not ready

## Security challenges are rising in relevance

Only a minority of organizations appear prepared to address the security challenges introduced by AI. As presented in previous sections, security ranks first among capability gaps, with 57% facing a deficit in AI security management, and second among understaffed technical domains, with 40% understaffed (Figure 8).

These numbers point to a significant vulnerability as AI systems become more capable and more deeply embedded in organizational operations. Agentic AI amplifies the risks even further. A non-deterministic system that can read and write files, query databases, trigger workflows, or call external APIs introduces risks that traditional application security models were not designed to handle and that require security expertise already missing in organizations.

Organizations are increasingly aware of the importance of security. As observed in Figure 9, security and privacy concerns ranked last in 2024 at 17%, climbed to second place in 2025 at 45%, and topped the list in 2026 at 48%.

**FIGURE 8**

### Organizations report substantial talent and capability gaps in security

LEFT: 2026 Tech Talent, Q19, Sample Size = 400, DKNS excluded

RIGHT: 2026 Tech Talent, Q25, Sample Size = 400, excluding Not Applicable, Gap = % reporting significant or minor capability gap

**40%**

of organizations report they are understaffed in

**cybersecurity and compliance**

#2 among understaffed tech domains

**57%**

of organizations report capacity gap in

**AI security & risk management**

#1 among technical capability gaps

## Security concerns are the leading barrier to getting value from AI

When asked specifically what prevents their organizations from getting value from AI, respondents rank security concerns first at 43% (Figure 10). Security concerns rank ahead of cost management challenges (36%), general skills gap (34%), legacy systems limitations (30%), insufficient business and domain knowledge (26%), and lack of infrastructure skills (21%).

Unlike deterministic software, where the same input produces the same output, systems based on generative AI behave probabilistically, and the same input can yield different actions across runs, making their behavior difficult to anticipate, test, or control.

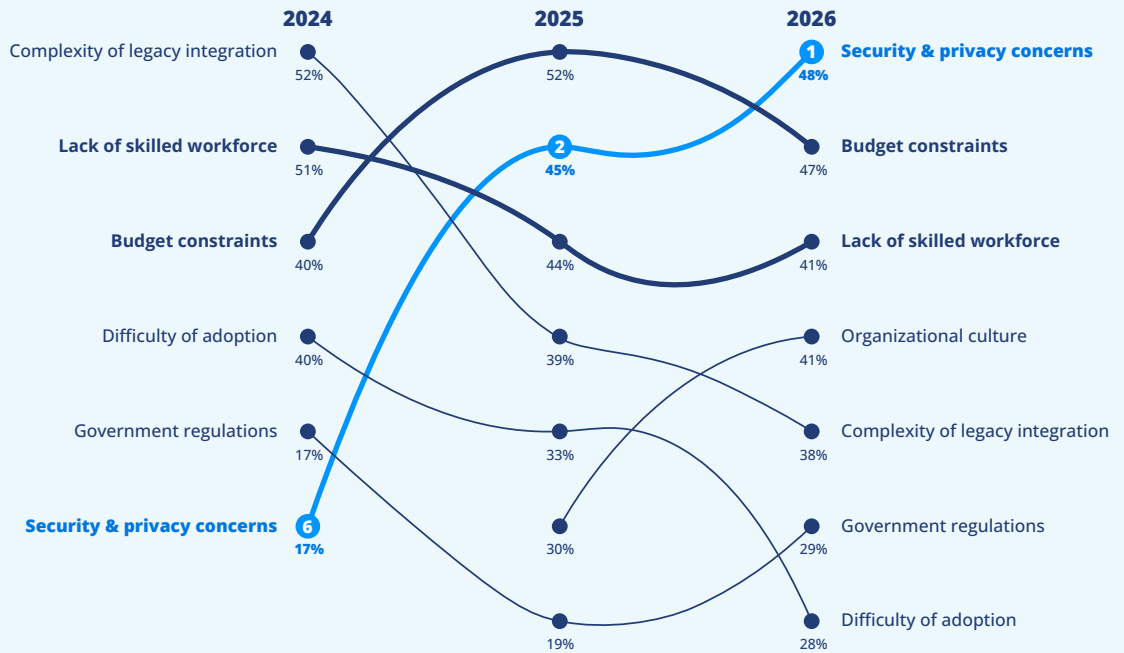
AI agents can operate autonomously by sending emails, querying databases, calling APIs, and making decisions. These activities expand the attack surface and create attack vectors that most organizations are not yet prepared to address, such as prompt injection, supply chain vulnerabilities in tools, and agents operating across trust boundaries without human oversight.

The security community has begun formalizing responses to these risks. OWASP (the Open Worldwide Application Security Project) launched its Top 10 for LLM Applications as a community-driven effort to identify and address security issues specific to AI applications<sup>22</sup>. In May 2025, the U.S. National Security Agency (NSA), alongside other federal agencies,

**FIGURE 9: SECURITY AND PRIVACY CONCERNS BECAME THE TOP BARRIER TO ADOPTING NEW TECHNOLOGIES IN 2026**

What are the primary challenge(s) of adopting new technologies? (select all that apply)

2024 Tech Talent Q47 (N=77), 2025 Tech Talent Q18 (N=556), 2026 Tech Talent Q18 (N=400), DKNS excluded



### Security Rising Fast

Security & privacy concerns jumped from last place in 2024 to #1 in 2026 — the fastest-rising barrier.

#6 17%  
in 2024

#2 45%  
in 2025

#1 48%  
in 2026


22 OWASP Gen AI Security Project, "OWASP Top 10 for LLM Applications 2025," Nov. 17, 2024. <https://genai.owasp.org/resource/owasp-top-10-for-llm-applications-2025/>



**FIGURE 10: TOP BARRIERS TO REALIZING VALUE FROM AI**

### What are your organization's top barriers to getting value from AI? (select up to three)

2026 Tech Talent, Q26, Sample Size = 400, Total Mentions = 805, DKNS excluded (2%), options with low response rates are omitted for clarity



released joint guidance titled “AI Data Security: Best Practices for Securing Data Used to Train & Operate AI Systems,” outlining best practices for securing data across the AI system lifecycle, including the data supply chain and protection against malicious or unauthorized modification.<sup>23</sup>

23 National Security Agency, “NSA’s AISC Releases Joint Guidance on the Risks and Best Practices in AI Data Security,” May 22, 2025. <https://www.nsa.gov/Press-Room/Press-Releases-Statements/Press-Release-View/Article/4192332/nsas-aisc-releases-joint-guidance-on-the-risks-and-best-practices-in-ai-data-se>

# The value of institutional knowledge

## Upskilling is the preferred strategy to address the lack of a skilled workforce

Workforce skills remain a major barrier to technology adoption. As shown in Figures 9 and 11, the lack of a skilled workforce has been among the top 3 barriers to adopting new technologies since 2024.

Organizations have a clear preference for how to address the technical talent gaps. Rather than relying on external hiring, most organizations are looking inward. Figure 12 shows upskilling existing staff (57%) as the primary response to talent gaps, ahead of hiring new technical staff (49%) and upskilling inexperienced staff (44%).


Existing staff carry institutional knowledge of systems, processes, codebases, and organizational context that cannot be hired in from outside. Building new technical competencies on top of the existing institutional knowledge foundation is less disruptive than onboarding external talent, who must first learn the environment before contributing. Not surprisingly, 94% of organizations consider upskilling important, very important, or extremely important (Figure 13).

“Organizations can’t buy institutional knowledge.”

**FIGURE 11: LACK OF SKILLED WORKFORCE IS AMONG THE TOP 3 BARRIERS TO ADOPTING NEW TECHNOLOGIES**

What are the primary challenges of adopting new technologies? (select all that apply)

2026 Tech Talent, Q26, Sample Size = 400, Total Mentions = 805, DKNS excluded (2%), options with low response rates are omitted for clarity



**FIGURE 12: UPSKILLING/CROSS-SKILLING IS THE TOP STRATEGY TO ADDRESS TECHNICAL TALENT GAPS**

### How organizations address talent gaps

2026 Tech Talent, Q27, Sample Size = 400, Total Mentions = 812, DKNS excluded (3%)



**FIGURE 13: UPSKILLING IS VIEWED AS IMPORTANT BY NEARLY ALL ORGANIZATIONS**

### How important is upskilling?

2026 Tech Talent, Q28, Sample Size = 400, DKNS excluded (1%)



## Upskilling is preferred across the technology stack

The same preference for upskilling extends to implementing AI initiatives. Figure 14 shows that upskilling the existing workforce in AI (50%) is the top strategy in organizations' plans to address their core AI activities, ahead of leveraging open source frameworks, models, and tools (43%), using AI-as-a-service platforms (36%), and staffing up and building in-house (33%).

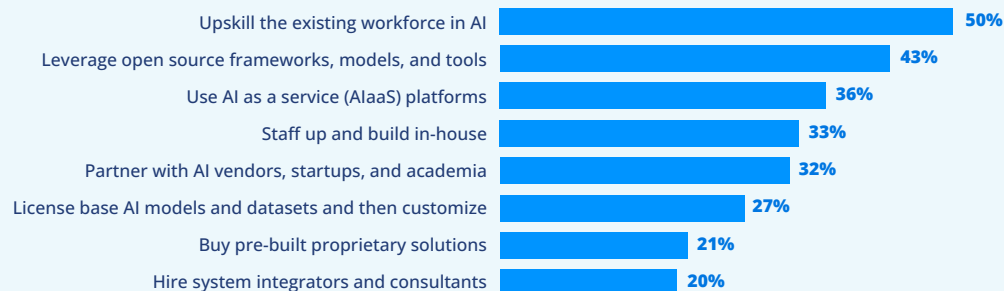
This finding suggests that effectively deploying AI is not purely a technical challenge. It also requires understanding the business processes, data flows, organizational context, and domain-specific nuances that



**FIGURE 14: UPSKILLING THE EXISTING WORKFORCE IS THE TOP STRATEGY FOR ADDRESSING AI CORE ACTIVITIES**

How does your organization plan to address its AI core activities? (select all that apply)

2026 Tech Talent, Q24, Sample Size = 400, Total Mentions = 1,043, DKNS excluded (4%), options with low response rates are omitted for clarity




shape how AI systems should behave and where they are likely to fail. That knowledge is largely tacit. It lives in the people who have worked within the organization, and it is not easily documented, transferred, or replicated by external hires. An AI system implemented by people who lack that context may be technically sound but operationally misaligned. Upskilling existing staff in AI, rather than bringing in AI specialists from outside, keeps that institutional knowledge in the loop.

Upskilling outpaces hiring new staff, not only in AI but across all surveyed technical domains (Figure 15). The margins are widest in cloud and containers (67% upskill vs. 16% hire), DevOps and site reliability (67% upskill vs. 16% hire), and AI/ML (51% upskill vs. 16% hire).

**FIGURE 15**

Upskilling beats hiring across priority domains

2026 Tech Talent, Q40, Sample Size = 400, DKNS excluded (4% to 16%). Not a priority is also excluded (6% to 48%). 2025 Tech Talent, Q42, Sample Size = 556, DKNS excluded (5% to 16%).<sup>24</sup>



24 Sorted by priority areas. Only the top 12 are shown. The average ratio is calculated for all domains, not only those shown in the figure. The ratio is calculated by dividing the percentage of upskill/cross by hire new technical staff, detailed data available in Appendix A6, including "Not a priority".

(63% vs. 17%), and platform engineering (60% vs. 15%), precisely the infrastructure domains critical for AI deployment. Cybersecurity, which faces the largest capability gap (Figure 6), shows the strongest preference for hiring new staff at 24%, suggesting that some organizations may perceive the security skills deficit as too specialized to close through internal development alone. Even so, upskilling still leads there at 53%.

The preference for upskilling is increasing over time. In 2026, organizations are 3.5 times more likely to upskill than to hire across all surveyed domains, up from 3.1 times the previous year. For organizations, this means that closing the skills gaps depends increasingly on the effectiveness of their internal learning programs.

## Upskilling vs. hiring: benefits and challenges

**FIGURE 16**

### Upskilling and hiring offer distinct tradeoffs in workforce development

2026 Tech Talent Survey, Q37 (n = 400). Top 3 responses displayed (labels abbreviated); complete response distributions are provided in Appendices A7-A10.



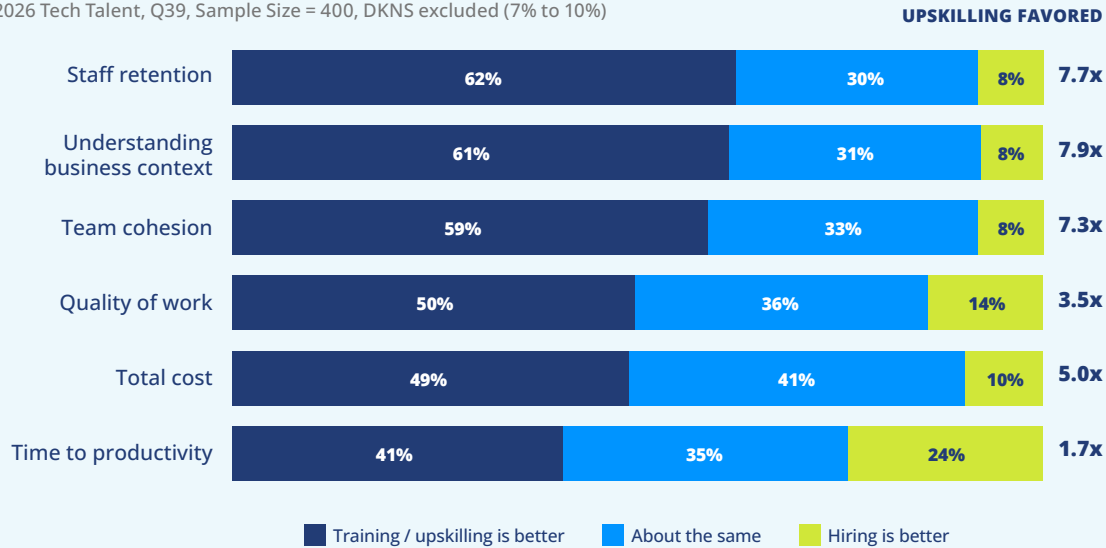
The preference for upskilling is widespread but not without trade-offs (Figure 16). The top benefits of upskilling relate to career development and job satisfaction (51%), the ability to redeploy staff more effectively than new hires (46%), and a growth pathway for junior talent (45%). Challenges also exist: 42% of organizations believe that building a continuous learning culture takes effort and time, 37% find upskilling can be ineffective in highly specialized roles, and 33% believe it diverts resources from other priorities.

External hiring also comes with benefits and challenges. External candidates bring fresh perspectives and innovative approaches (52%), arrive with the required skills (43%), and allow specific gaps to be targeted with precision (43%), making it attractive when the required expertise is too specialized to develop internally in a reasonable timeframe. The costs, however, are substantial: 46% consider that it can delay projects, 40% report that recruitment is frequently costly and fails to identify the right candidate, and 35% consider that verifying claimed technical skills is difficult.

**FIGURE 17: UPSKILLING OUTPERFORMS HIRING ACROSS KEY DIMENSIONS**

How do these approaches compare in your experience?

2026 Tech Talent, Q39, Sample Size = 400, DKNS excluded (7% to 10%)



When organizations compare the two approaches side by side (Figure 17), upskilling outperforms hiring across all surveyed dimensions. The preference for upskilling is especially pronounced in dimensions strongly tied to accumulated organizational knowledge and social capital, such as understanding of business context (7.9x advantage for upskilling), staff retention (7.7x), and team cohesion (7.3x). Other dimensions with clear preference for upskilling include total cost (5x) and quality of work (3.5x). The narrowest margin is time to productivity at 1.7x.

Our survey confirms that the hiring cycle is much slower than upskilling and can be risky (Figure 18). Hiring and onboarding new staff takes 53% longer than upskilling, and 28% of new hires leave within six months (up from 19% the previous year). These findings reinforce the idea that business context, tacit knowledge, and institutional familiarity can't be taught quickly. Additionally, when nearly one in three new employees leave within half a year, onboarding effort is lost and organizations must restart the process.



**FIGURE 18**

Hiring is slower and riskier than upskilling

2026 Tech Talent, Q30, Q31, Q32, Q36, segregated by Q8 and Q13, Sample Size = 400, weighted averages, detailed data available in Appendices A11-A14

Hiring and onboarding new staff takes **53%** longer than upskilling (vs. 62% in 2025)

**28%** of new employees leave within 6 months (vs. 19% in 2025)




Technical training is among the top strategies to retain technical talent

Retention is one of the main benefits organizations associate with an upskilling strategy (Figure 19). For retaining technical talent, technical growth opportunities (94%) rank first, tied with work-environment factors such as remote work and flexible hours. Technical training, including training and certification programs, ranks third at 93%, just one percentage point behind the leaders and

**FIGURE 19: STRATEGIES TO RETAIN TECHNICAL TALENT**

What strategies does your organization employ to retain technical talent?  
Rate the effectiveness of each:

2026 Tech Talent, Q29, Sample Size = 400, excluding DKNS (3% to 7%) and Not Offered (9% to 20%)



ahead of compensation (91%) and career growth opportunities (91%).

These results suggest that technical professionals weigh learning and development at least as heavily as compensation when deciding whether to stay. Open source culture, including dedicated time for open source contributions and a clear OSS policy, ranks last among the surveyed factors. Even so, 81% of respondents recognize its effectiveness.

These findings reinforce a central theme of this report: technical talent is more likely to stay where it has room to grow. Organizations that invest in continuous development are not only closing skills gaps but also creating conditions that encourage skilled professionals to stay.

“Technical talent is more likely to stay where it has room to grow.”

The factors organizations value most in technical assessments favor internal candidates

When assessing candidates’ skills for a technical position, hands-on experience from prior employment (93%) ranks first, followed by portfolio and examples of IT project work (87%) (Figure 20). Both are easier to evaluate for internal candidates than for external hires. Organizations have direct

**FIGURE 20: IMPORTANCE OF FACTORS USED TO ASSESS CANDIDATES’ TECHNICAL SKILLS**

What employers value most

2026 Tech Talent, Q33, Sample Size = 400, DKNS excluded (1% to 3%), aggregation of extremely important, very important, and important, detailed data available in Appendix A15



visibility into existing employees' track records, including how they have performed under real conditions, which systems they have worked on, and how they handle unfamiliar technical challenges.

That evidence is difficult to infer from a resume or reconstruct through an interview process. When the primary criteria for technical assessment are experience and demonstrated work, internal development has a clear advantage.

Certifications of skills (76%) and formal college or university degrees (74%) trail the top two criteria but remain broadly valued. Consistent with the last two years' surveys, certifications ranked higher than formal college or university degrees. Certifications are more narrowly focused and more directly tied to the specific technical domains organizations need. For organizations running upskilling programs, certifications provide a structured path for closing defined gaps while also carrying weight in internal assessments and retention.

## Conclusion

Findings from the 2026 State of Tech Talent Report suggest that the primary challenge facing the talent market is a skills crisis driven by technological demands, rather than a lack of job opportunities. According to survey respondents, AI is acting as a net driver of job growth in IT, with a positive net hiring effect of +31% expected for 2026. However, this growth is not uniform. Large enterprises are the only segment reporting a negative net hiring effect, suggesting that recent layoffs may be concentrated among the largest firms while smaller organizations absorb more of the demand.

The report points to a full-stack readiness problem, with organizations needing capabilities that extend beyond AI specialization alone. While AI/ML engineering is notably understaffed (47%), capability gaps are also extensive across strategic domains, including security and risk management and AI operations and monitoring. Security concerns have become the leading barrier to adopting new technologies. These results align with the low implementation rates of foundational AI infrastructure such as the PARK stack.

To bridge these talent gaps, organizations largely prefer internal development. Upskilling existing staff is favored over external hiring, offering important advantages in preserving institutional knowledge, understanding the business context, and achieving faster time-to-productivity and higher retention rates. These findings suggest that technical professionals are more likely to stay in organizations where they have opportunities to grow, with technical training and growth opportunities ranking as highly as compensation in retention decisions.

Based on survey results, recommendations for hiring managers include:

1. Prioritize internal talent development and upskilling to reap the benefits of an internal candidate while reducing the costs and risks of external hiring.
2. Foster a continuous learning environment to increase the effectiveness of upskilling and boost retention for employees seeking technical training and growth opportunities.
3. Emphasize practical experience and certifications in assessment and actively support certifications as a structured path to close skills gaps.

For those seeking software roles, recommendations include:

1. Focus on full-stack skills to support the operationalization of AI, including cybersecurity and risk management, AI operations and monitoring, and platform engineering.
2. Build a strong portfolio and pursue certifications to document and demonstrate practical abilities.



# Methodology and demographics

## About the survey

This study is based on an online survey conducted by Linux Foundation Research in February 2026. The survey aimed to understand trends in technical hiring and the effects of AI-related disruption on the workforce. We broadly advertised the survey to Linux Foundation subscribers, members, partner communities, and social media. To mitigate sampling biases, we also hired a panel provider. We addressed data quality through extensive prescreening, survey screening questions, consistency checks, and data quality review to ensure that respondents had sufficient professional experience to answer questions accurately on behalf of their organization. We only considered complete responses to the survey. After the data quality filtering, our final sample comprised 400 valid responses.

The survey comprised 40 questions covering screening, respondent demographics, AI's impact on organizations, and approaches to addressing talent management needs. The dataset driving the analysis in this report and survey frequencies can be found on Data.World (see below).

The target audience included respondents who met the following criteria:

- Must be responsible for hiring, recruiting, or training information technology (IT) professionals
- Must pass an attention check question.
- Must be currently employed by an organization.

A total of 1,077 respondents began the survey and 530 completed it. After data-quality screening, the final analyzed dataset comprised 400 responses. The margin of error for this sample size is +/- 4.1% at a 90% confidence level and +/- 4.9% at a 95% confidence level. The data was primarily segmented by region, organization size, and industry.

Although respondents were required to answer nearly all questions in the survey, a provision was made when a respondent was unable to answer a question. This was accomplished by adding a "Don't know or not sure" (DKNS) response to the list of responses for every question. However, this creates a variety of analytical challenges. One approach was to treat a DKNS just like any other response, so that the percentage of respondents who answered the DKNS is known. This approach has the advantage of showing the exact

distribution of the data collected. The challenge with this approach is that it can distort the distribution of valid responses, i.e., responses where respondents could answer the question. Therefore, most of the analyses in this report exclude DKNS responses. This is because we can classify the missing data as either missing at random or missing completely at random. Excluding DKNS data from a question does not change the distribution of data (counts) for the other responses, but it does change the size of the denominator used to calculate the percentage of responses across the remaining responses. This has the effect of proportionally increasing the percentage values of the remaining responses. Where we have elected to exclude DKNS data, the footnote for the figure indicates “DKNS responses excluded from the analysis.”

The percentage values in this report may not total exactly 100% due to rounding.

## Data.World access

LF Research makes each of its empirical project datasets available on Data.World ([data.world/thelinuxfoundation](https://data.world/thelinuxfoundation)). Included in this dataset are the survey instrument, raw survey data, screening and filtering criteria, and frequency charts for each question in the survey. Access to Linux Foundation datasets is free but does require you to create a data.world account.

## Respondent demographics

Figure 21 presents the respondent demographics. Most respondents (66%) are hiring managers for technical teams. The vast majority (84%) are employed full-time. Professionally, they are most heavily concentrated in software development (28%) and systems operations (21%), with C-level executives making up 6% of respondents. These professionals work predominantly for European (38%) and Asia-Pacific (29%) organizations, followed by US and Canada (22%). They are spread across organizations of varying sizes, with 27% working in smaller companies (1–249 employees), 40% in mid-sized companies (250–4,999 employees), and 32% in large enterprises (5,000+ employees). Most respondents work for technology-focused companies, with 52% working for organizations whose primary revenue comes from IT products or services, while 40% work for companies providing industry-specific products or services.

Some demographics have been regrouped to facilitate a more insightful analysis. For the original source data and study frequencies, please see the data.world dataset and access as described above.

## FIGURE 21: RESPONDENT DEMOGRAPHICS


### What best describes your role in training or hiring IT professionals?

2026 Tech Talent, Q2, n=400




### Professionally, which role or field do you most closely identify with?

2026 Tech Talent, Q7, n=400




### Which option best describes the company or entity you work for?

2026 Tech Talent, Q9, n=400




### What is your current employment status?

2026 Tech Talent, Q5, n=400




### In what country or region does your organization have its headquarters?

2026 Tech Talent, Q8, n=400



### Please estimate how many total employees are in the company or entity you work for.

2026 Tech Talent, Q13, n=400



# Appendix

## Appendix A1 2026 Tech Talent, Q21 by Q8, by Q9, by Q13, data for 2025, Sample Size = 400, DKNS and Not Applicable excluded

HOW HAS AI IMPACTED OR WILL IMPACT THE HEADCOUNT OF YOUR ORGANIZATION? (Q21)	DECREASED OUR HEADCOUNT	NO IMPACT TO HEADCOUNT	INCREASED OUR HEADCOUNT
1-249 employees	3%	66%	31%
250-4,999 employees	16%	38%	45%
5,000-19,999 employees	10%	38%	52%
20,000+ employees	28%	48%	24%
Asia-Pacific	16%	34%	50%
Europe	11%	55%	34%
US or Canada	18%	46%	36%
IT Industry	17%	46%	38%
End-user organization	12%	44%	44%

## Appendix A2 2026 Tech Talent, Q22, Sample Size = 400, not applicable and DKNS excluded (11% to 15%)

HOW HAS AI IMPACTED YOUR TECHNICAL WORKFORCE IN THE FOLLOWING AREAS DURING 2025? (Q22)	INCREASE	NO CHANGE	DECREASE
AI-specific roles	63%	33%	3%
Software development positions	40%	47%	12%
Technical management roles	30%	63%	8%
IT operations staff	28%	60%	12%
QA/testing positions	32%	51%	16%
Entry-level technical positions	29%	49%	22%

## Appendix A3 2026 Tech Talent, Q19, Sample Size = 400, DKNS excluded (2% to 14%)

RATE YOUR ORGANIZATION'S CURRENT STAFFING CAPACITY IN THESE STRATEGIC DOMAINS: (Q19)	SIGNIFICANTLY UNDERSTAFFED	SOMEWHAT UNDERSTAFFED	ADEQUATELY STAFFED	WELL-STAFFED	OVER-STAFFED
AI & ML engineering and operations	17%	30%	24%	22%	7%
Cybersecurity and compliance	11%	28%	35%	22%	3%
FinOps and cost optimization	13%	23%	35%	23%	5%
Platform engineering	9%	26%	33%	28%	5%
Cloud computing (public and private)	6%	22%	39%	28%	5%

## Appendix A4 2026 Tech Talent, Q25, Sample Size = 400, excluding Not Applicable (7% to 10%)

RATE YOUR ORGANIZATION'S CURRENT CAPABILITY IN THESE AREAS FOR AI SUCCESS: (Q25)	SIGNIFICANT GAP	MINOR GAP	ADEQUATE	STRONG CAPABILITY
AI security and risk management	24%	34%	26%	17%
AI operations and monitoring	20%	38%	28%	15%
Cost optimization for AI workloads	23%	31%	31%	15%
Expertise to build/manage AI infrastructure	18%	27%	38%	17%
Change management and communication	17%	27%	41%	16%
Data management	14%	26%	40%	20%
Infrastructure and platform capabilities (Kubernetes, cloud-native systems)	16%	20%	40%	25%
Business domain expertise	12%	22%	41%	24%

## Appendix A5 2026 Tech Talent, Q20, Sample Size = 400, DKNS excluded (3% to 15%)

WHAT PROGRESS HAS YOUR ORGANIZATION MADE IN IMPLEMENTING EACH LAYER OF THE PARK STACK? (Q20)	NOT APPLICABLE	NOT STARTED	PLANNING	IN PROGRESS	PARTIALLY DEPLOYED	FULLY DEPLOYED
High-level framework such as PyTorch for model development or refinement	19%	8%	16%	26%	20%	10%
AI foundation models that can be tuned, customized, and deployed	8%	14%	23%	33%	16%	5%
Distributed compute engineering such as Ray for scaling applications and handling systems challenges	25%	10%	23%	17%	12%	13%
Container orchestration such as Kubernetes for provisioning resources	5%	6%	20%	25%	41%	4%

## Appendix A6 2026 Tech Talent, Q40, Sample Size = 400, DKNS excluded (4% to 16%). Sorted by priority areas.

These numbers may be different than the ones presented in the text, which were recalculated excluding the “not a priority” responses.

FOR THE FOLLOWING TECHNOLOGY AREAS, WHICH APPROACH WOULD YOU PRIORITIZE IN 2026 TO MEET YOUR ORGANIZATION'S NEEDS? (Q40)	UPSKILL OR CROSS-SKILL EXISTING TECHNICAL STAFF	HIRE NEW TECHNICAL STAFF	HIRE CONSULTANTS	NOT A PRIORITY
Cybersecurity	50%	23%	21%	6%
DevOps, CI/CD and site reliability	57%	15%	19%	8%
Cloud and containers	61%	15%	15%	10%
AI/ML	46%	19%	25%	10%
Data and analytics	51%	13%	22%	15%
Web & Application Development	44%	17%	24%	15%
Platform engineering	50%	12%	21%	16%
System engineering	49%	13%	20%	18%
System administration	47%	11%	20%	21%
Networking and edge	45%	9%	21%	25%
Open source and compliance best practices	39%	7%	24%	30%
Linux kernel and operating systems	40%	10%	19%	31%
Supply chain	32%	9%	26%	33%
Safety-critical systems	35%	13%	20%	33%
IoT and embedded	31%	9%	21%	39%
Open Source Program Offices (OSPO)	30%	7%	21%	43%
Visual effects, augmented/virtual reality	27%	10%	19%	45%
Open hardware	24%	7%	24%	45%
Quantum computing	21%	12%	22%	46%
Blockchain and decentralized identity	20%	7%	26%	48%

## Appendix A7

2026 Tech Talent, Q37, Sample Size = 400, Total Mentions = 1,148, DKNS excluded (2%)

WHAT ARE THE MAIN BENEFITS TO UPSKILLING OR CROSS-SKILLING STAFF? (SELECT ALL THAT APPLY) (Q37)	%
Presents career development opportunities to staff members, potentially leading to higher job satisfaction	51%
Produces employees with varied skillsets who can be redeployed more effectively than new hires	46%
Provides an ideal pathway for junior technical staff to expand their capabilities	45%
Brings people into roles who already possess familiarity with and have relationships with the company	45%
A more cost-effective approach to increasing technical expertise compared to hiring new employees	41%
Serves as an effective strategy for filling senior positions when technical headcount is difficult to find	39%
Works best when learning narrowly focused skills	21%
Other (please specify)	2%
Our organization sees no benefits to upskilling	1%

## Appendix A8

2026 Tech Talent, Q38, Sample Size = 400, Total Mentions = 865, DKNS excluded (4%)

WHAT ARE THE MAIN CHALLENGES TO UPSKILLING OR CROSS-SKILLING STAFF? (SELECT ALL THAT APPLY) (Q38)	%
Creating and nurturing an environment for continuous learning takes effort and time	42%
Upskilling takes time and can be ineffective at training for complex roles	37%
Resources are pulled away from other priority areas	33%
Finding the appropriate training materials is a challenge	31%
It can be difficult to translate new knowledge into practical applications	27%
Upskilling doesn't work for senior roles because you can't teach broad subject matter expertise	23%
Upskilling staff doesn't help us to fill new positions because we still have to backfill for the positions vacated	22%
Our organization has not experienced any challenges to upskilling	2%
Other (please specify)	2%

## Appendix A9

2026 Tech Talent, Q34, Sample Size = 400, Total Mentions = 1,045, DKNS excluded (2%)

WHAT ARE THE MAIN BENEFITS OF HIRING TECHNICAL STAFF? (SELECT ALL THAT APPLY) (Q34)	%
New technical employees can bring fresh perspectives and innovative approaches	52%
The candidate already has the skills required for the role	43%
Skill gaps can be addressed with precision	43%
Hiring individuals with multiple skills can cover multiple roles	42%
Access to a wider talent pool	30%
Increasing overall staff headcount means additional, instead of reallocated resources	27%
Little or no investment needed in their technical training	25%
Our organization sees no benefits to hiring	1%
Other (please specify)	1%

## Appendix A10

2026 Tech Talent, Q35, Sample Size = 400, Total Mentions = 968, DKNS excluded (3%)

WHAT ARE THE MAIN CHALLENGES TO HIRING FOR TECHNICAL STAFF? (SELECT ALL THAT APPLY) (Q35)	%
Problems in finding the candidate with the right skills can delay projects	46%
Recruitment is costly, time consuming, and often does not lead to the right candidate for the position	40%
Verifying claimed technical skills	35%
Not all tech talent is the same – it's challenging that they do not have a common benchmark of skills, methodologies	34%
Onboarding takes valuable internal resources away from other critical projects and is time consuming	31%
Hiring the wrong candidate and then having to onboard again is an issue	30%
There is not a streamlined way to recruit new technical staff with proven skills	25%
Our organization has not experienced any challenges to hiring	3%
Other (please specify)	3%

## Appendix A11

2026 Tech Talent, Q30, Sample Size = 400

ON AVERAGE, HOW LONG DOES IT TAKE TO HIRE A HEADCOUNT TO FILL AN OPEN TECHNICAL POSITION IN YOUR ORGANIZATION? (SELECT ONE) (Q30)	%
Less than a month	5%
1 to 3 months	33%
4 to 6 months	29%
7 to 9 months	17%
10 to 12 months	8%
More than 12 months	2%
Don't know or not sure	6%

## Appendix A12

2026 Tech Talent, Q31, Sample Size = 400

HOW LONG DOES THE ONBOARDING PROCESS TAKE FOR A NEW TECHNICAL HEADCOUNT TO REACH NORMAL PRODUCTIVITY? (SELECT ONE) (Q31)	%
Up to 1 month	13%
1 to 3 months	37%
4 to 6 months	36%
7 to 12 months	10%
More than 12 months	1%
Don't know or not sure	4%

## Appendix A13

2026 Tech Talent, Q36, Sample Size = 400

ON AVERAGE, HOW LONG DOES IT TAKE TO UPSKILL OR CROSS-SKILL AN EXISTING EMPLOYEE TO REACH NORMAL PRODUCTIVITY IN A NEW TECHNICAL ROLE? (SELECT ONE) (Q36)	%
Less than a month	2%
1 to 3 months	26%
4 to 6 months	38%
7 to 9 months	19%
10 to 12 months	7%
More than 12 months	4%
Don't know or not sure	5%

## Appendix A14

2026 Tech Talent, Q32, Sample Size = 400

ON AVERAGE, WHAT PERCENTAGE OF NEW TECHNICAL STAFF HIRES RESIGN OR WERE ASKED TO LEAVE WITHIN 6 MONTHS OF BEING ONBOARDED? (SELECT ONE) (Q32)	%
0-20%	45%
21-40%	18%
41-60%	19%
61-80%	8%
81-100%	1%
Don't know or not sure	10%

## Appendix A15 2026 Tech Talent, Q33, Sample Size = 400, DKNS excluded (1% to 3%)

HOW IMPORTANT ARE THE FOLLOWING EDUCATION AND EXPERIENCE FACTORS WHEN ASSESSING A CANDIDATE'S TECHNICAL SKILLS? (Q33)	EXTREMELY IMPORTANT	VERY IMPORTANT	IMPORTANT	NOT VERY IMPORTANT	NOT IMPORTANT AT ALL
Relevant hands-on experience based on previous employment	27%	46%	19%	6%	1%
A portfolio and examples of previous IT project accomplishments	14%	38%	34%	12%	2%
Certification of skills	7%	24%	45%	19%	5%
Formal college or university degree	10%	19%	45%	19%	7%

# Acknowledgments

Thanks to our sponsor KodeKloud for their participation in distributing the survey and sharing their insight. We also thank all the people who participated in the survey. Special thanks to Linux Foundation colleagues and community members for their involvement in the various stages of the research process: Hilary Carter, Clyde Seepersad, Mary Simpkins, Scott Punk, Mieko Sato, Mia Chaszeyka, and the LF Creative Services team.



KodeKloud

# About the authors

**MARCO GEROSA** is a full professor of Computer Science at Northern Arizona University and a research analyst at LF Research. His research on software engineering and open source software has resulted in over 300 publications in top-tier venues. Dr. Gerosa has a Ph.D., a master's in Informatics, and a B.S. in Computer Engineering. He is a Senior Member of the Institute of Electrical and Electronics Engineers (IEEE) and the Association for Computing Machinery (ACM). He supervised several Ph.D. and M.Sc. students who are now researchers in top institutions. He also has more than 20 years of teaching experience. For more information, visit [marcoagerosa.com](http://marcoagerosa.com).

**ADRIENN LAWSON** serves as Director of Quantitative Research at the Linux Foundation, where she leads data-driven initiatives to understand open source ecosystems. With expertise in social data science from the University of Oxford and a background spanning academic and governmental research, she brings methodological rigor to analyzing distributed collaboration networks. At the Linux Foundation, Adrienn leads a team conducting cross-sectional research across industry verticals and geographic regions to provide comprehensive insights into open source dynamics. Her work encompasses empirical investigations into regulatory compliance, the implications of AI, and sustainable funding models. She produces evidence-based recommendations that inform strategic decision-making within the open source community.

**ANNA HERMANSEN** is the Senior Manager, Researcher & Ecosystem for the Linux Foundation, where she supports end-to-end management of the Linux Foundation's research projects. She has conducted qualitative and systematic review research in health data infrastructure and the integration of new technologies to better support data sharing in healthcare, and has presented on this research work at conferences and working groups. Her interests lie at the intersection of health informatics, precision medicine, and data sharing. She is a generalist with experience in client services, program delivery, project management, and writing for academic, corporate, and web user audiences. Prior to the Linux Foundation, she worked for two different research programs, the Blockchain Research Institute and BC Cancer's Research Institute. She received her Master of Science in Public Health and a Bachelor of Arts in International Relations, both from the University of British Columbia.

 [x.com/linuxfoundation](https://x.com/linuxfoundation)

 [facebook.com/TheLinuxFoundation](https://facebook.com/TheLinuxFoundation)

 [linkedin.com/company/the-linux-foundation](https://linkedin.com/company/the-linux-foundation)

 [youtube.com/user/TheLinuxFoundation](https://youtube.com/user/TheLinuxFoundation)

 [github.com/LF-Engineering](https://github.com/LF-Engineering)

May 2026



Copyright © 2026 The Linux Foundation

This report is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International Public License.

To reference this work, please cite as follows: Marco Gerosa, Adrienn Lawson, and Anna Hermansen, "2026 State of Tech Talent: Not a jobs crisis, but a skills crisis with an upskilling answer," foreword by Clyde Seepersad, The Linux Foundation, May 2026.



Founded in 2021, [Linux Foundation Research](#) explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. By leveraging project databases and networks and committing to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations worldwide.



[Linux Foundation Education](#) builds the certifications and skills that power the technologies the world runs on. As the training and certification arm of the Linux Foundation, it is uniquely positioned at the source of innovation, working directly with the maintainers, projects, and communities shaping modern infrastructure. Its vendor-neutral certifications set the standard for proving real-world capability, while its training helps individuals build the skills to earn them. Together, they enable professionals to grow credible, portable careers and help organizations develop teams that can adopt, scale, and secure critical technologies across cloud, AI, and emerging domains.