Kyle & Co

INDEPENDENT MARKET RESEARCH

The State of

Candidate Fraud Detection *& Prevention*

How an HR problem became a security problem and what the market is building to meet it

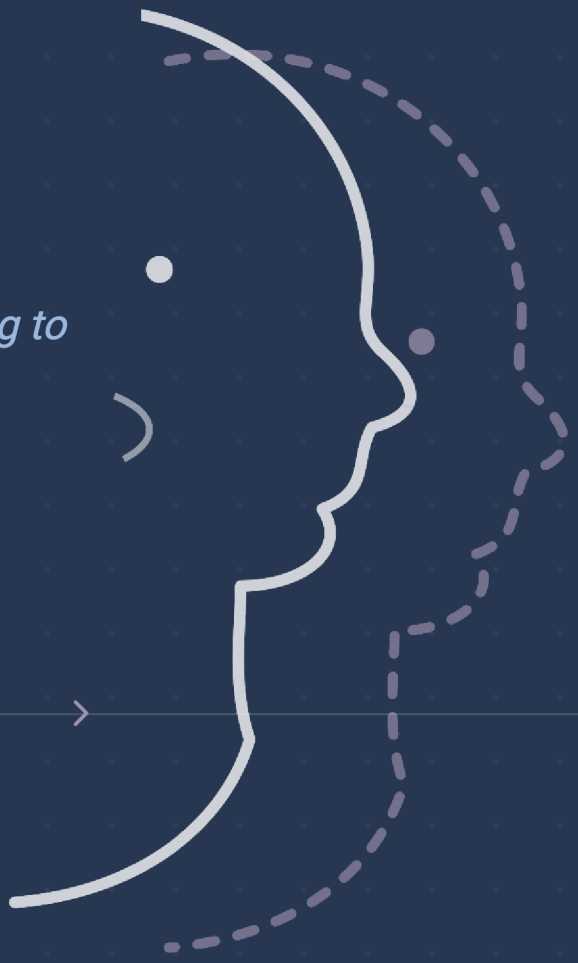

A RESEARCH REPORT BY KYLE & CO.

June 2026

CONTENTS

Contents

FRONT MATTER

THE REPORT

CLOSING

Introduction

Most organizations have encountered some form of candidate fraud. A candidate exaggerates responsibilities on a resume. Someone receives unauthorized assistance during an assessment. A reference cannot be verified.

Hiring teams have long managed varying degrees of misrepresentation, and for many years the consequences were relatively contained. The result was usually a poor hiring decision, additional recruiting costs, or a role that had to be reopened sooner than expected.

> ***What has changed is the scale, sophistication, and accessibility of the tools involved.***

Artificial intelligence (AI) can generate convincing resumes, support candidates during interviews, and help create synthetic professional identities. Remote hiring has expanded access to talent while reducing many of the traditional signals organizations once relied upon to establish trust and authenticity. At the same time, organizations are confronting more sophisticated forms of fraud, including identity manipulation, deepfake interviews, coordinated fraud networks, and workforce infiltration schemes designed to gain access to systems, data, and intellectual property.

As a result, candidate fraud is no longer solely a recruiting concern. It increasingly intersects with security, compliance, legal, operational, and business risk discussions.

Organizations are responding in a variety of ways. Some are strengthening hiring processes and verification practices. Others are introducing identity verification, interview integrity, and fraud detection technologies. Many are discovering that the greatest challenge is not technology at all, but establishing ownership, governance, and collaboration across Talent Acquisition (TA), HR, Security, Legal, and IT.

This report examines the current candidate fraud landscape, the technologies emerging to address it, and the organizational decisions required to build an effective response. It explores how fraud appears across the hiring lifecycle, where defenses remain weakest, how the market is evolving, and what organizations can do today to reduce risk.

Scope

This report is intended for Talent Acquisition leaders, Security leaders, HR executives, and other stakeholders responsible for protecting hiring integrity and managing organizational risk. It examines how candidate fraud is evolving, where vulnerabilities emerge throughout the hiring lifecycle, how the market is responding, and what organizations can do today to strengthen their defenses.

The report combines market analysis, practitioner insights, documented incidents, and implementation guidance to help organizations better understand the candidate fraud landscape and make more informed decisions about governance, processes, and technology investments.

It is not a vendor ranking, procurement scorecard, or product benchmark. Rather than identifying a single “best” solution, the report provides a framework for evaluating risks, understanding market capabilities, establishing ownership, and selecting approaches that align with an organization’s specific needs and risk profile.

ADDITIONALLY, A NOTE ON VENDOR CLAIMS

Additionally, a note on vendor claims: Most accuracy figures, detection rates, and fraud-blocked statistics in this market are self-reported by vendors against methodologies they designed, on populations they selected. Independent validation exists for a small subset of providers—primarily in identity verification and deepfake detection, where certification infrastructure from adjacent fields already existed. Throughout this report, we note where claims are independently validated and where they are not. Buyers should weight structural questions and third-party certifications more heavily than accuracy percentages when evaluating vendors.

Executive Summary

Candidate fraud increasingly affects hiring outcomes, operational risk, information security, compliance, and organizational reputation. As a result, many organizations are treating it as a cross-functional business risk rather than a recruiting issue.

The Fastest Facts

Five key points to set the stage

1 Candidate fraud is an organizational risk, not just a hiring risk.

Candidate fraud now extends beyond résumé inaccuracies and hiring quality concerns. Organizations are encountering AI-generated applications, identity fraud, deepfake interviews, assessment cheating, coordinated reference networks, and workforce infiltration attempts. These risks can affect hiring outcomes, operational effectiveness, information security, compliance, and organizational reputation.

2 Fraud occurs throughout the hiring lifecycle, and vulnerabilities differ by stage.

Different forms of fraud emerge at different points in the hiring process. Application intake, assessments, interviews, references, onboarding, and post-hire activities each present unique risks and require different controls. Organizations benefit from evaluating fraud through the entire hiring lifecycle rather than focusing on a single checkpoint or technology solution.

3 The market is growing rapidly, but no single approach has emerged as the standard.

The candidate fraud market continues to expand as vendors from recruiting, cybersecurity, identity verification, assessment, and workforce technology enter the space. While capabilities are improving quickly, categories continue to evolve, vendor approaches vary significantly, and evaluation frameworks are still emerging.

4

The greatest challenge is often ownership, not technology.

Talent Acquisition is frequently the first team to encounter candidate fraud, but fraud prevention increasingly requires collaboration among Security, Legal, Compliance, IT, and business leadership. Organizations making the most progress have established clear ownership, defined escalation procedures, and shared accountability across functions.

5

Technology supports the solution, but programs determine outcomes.

Successful organizations combine technology, process, governance, training, and cross-functional collaboration. They treat candidate fraud as an operational capability rather than a standalone technology purchase. Clear ownership, documented response procedures, recruiter enablement, and ongoing review processes consistently appeared in the strongest programs observed during this research.

The Deeper Findings

Six things taking the fraud conversation to the next level

The job of a research report is partly to organize the obvious. The more useful job is to surface what is not obvious—the findings that complicate the conventional reading of the market, that the vendor pitches may obscure, or that emerge only when enough sources are stacked next to each other for the patterns to show through. Six such findings recurred consistently enough across the research that you need to know.

1 The metrics that measure recruiting success cannot tell whether a fraud program is working.

Time-to-fill goes up. Cost-per-hire goes up. Pipeline conversion rates go down. These are the metrics a TA organization is judged on—and they are, mechanically, the same signals you would see from a TA team running a serious fraud program and from a TA team that has simply gotten worse.

This is more than a measurement problem. It is a structural disincentive: a recruiter or TA leader who voluntarily adds friction to catch fraud is, by every standard scorecard, hurting their own performance. The organizations doing this work seriously and the organizations failing at it look identical from the metrics dashboard — and may, in the most uncomfortable version of the finding, be the same organization. The team catching the most fraud may have the worst time-to-fill, not despite their efforts but because of them.

Until TA scorecards begin to include some measure of hiring integrity — quality-of-hire over time, post-hire termination rates from misrepresentation, audit results — the incentive structure quietly punishes the right behavior. None of the practitioners interviewed for this research had a TA scorecard that accounted for fraud prevention at all.

---

2 The embedded paradox: embedded fraud detection works best for the organizations that least need it.

The intuitive read of “embedded fraud detection in your ATS” is that it democratizes the capability—fraud screening, for the masses, included in a tool the team already uses. The research suggests something closer to the opposite.

The embedded model produces signals. Signals require a human to adjudicate them. Adjudication requires someone with the time, the mandate, and the training to make a defensible call when a candidate is flagged. A large organization with a security team, a TA operations function, and an established cross-functional structure can absorb that adjudication workload. A two-person TA team running on time-to-fill cannot.

One of the product leads we interviewed said this directly: the company has “no wonderful answer” for the small organization with no security partner to escalate signals to. The feature works. The organization around it does not.

The paradox is that the organizations least equipped to act on a flag are the ones most likely to be sold an embedded feature as their fraud strategy—because the embedded layer is also the easiest to adopt. Capability arrives at the desks of the people least positioned to use it. Embedded is fine; embedded without an adjudicator is a feature pretending to be a program.

3 Agreement across the architecture divide is the most reliable finding in the research.

Vendors often disagree with each other on root problems and ideal solutions. So when four vendors who compete with each other, who are built on entirely different architectures, who come from different categories all independently arrive at the same design choices, that convergence is especially remarkable. Four of the vendors interviewed or reviewed for this report—Ashby (an embedded ATS), Socure (a standalone identity platform), CodeSignal (an assessment platform), and Findem (a talent-data platform building an authenticity layer)—all landed independently on the same three commitments:

- **Surface signals rather than verdicts;**
- **Explain every signal so the recruiter can interrogate it;**
- **And refuse automated rejection—the human makes the decision.**

Four vendors from four different categories converging on the same answer is not a coincidence. It is a structural finding about what the market has already learned—that black-box scores produce trust problems, that automated employment decisions produce legal exposure, and that the only design that survives contact with regulation, recruiters, and reality is the one that keeps a human in the loop.

Buyers should be wary of vendors who have not yet arrived at this conclusion. The four that have got there based on learnings the market is steadily ratifying.

---

4 The highest-value fraud signal in the market is the least-used.

When we asked vendors in this space about the most sophisticated detection capabilities, we heard about biometric liveness, behavioral telemetry, voice deepfake classifiers, and identity-graph reasoning. But the single most reliable, least gameable, lowest-friction fraud signal in this entire landscape is much less exciting: the IP address and device fingerprint of the references a candidate submits.

If five "references" submit their feedback from the same IP address, or the same device, that is not a probability—that is the signature of a single person filling out their own reference checks, or of a small ring coordinating to do it. It requires no biometric data, raises no compliance issues, and is essentially impossible to fabricate at scale because the infrastructure cost of getting five distinct devices in five distinct geographies for one fraudulent application is not worth the payoff.

Most organizations are not able to monitor these signals today. Either their reference vendor does not surface it, or—more commonly—they have not automated reference checks in the first place. The most effective check in the funnel is the one most teams have stopped doing.

5 Several of the fraud vendors are themselves being attacked.

A pattern that emerged across multiple interviews rarely surfaced publicly: the vendors building fraud detection products are also frequent targets of the fraud they detect. Socure reports that more than a quarter of its own applicants are fake, and even caught a sophisticated DPRK candidate it nicknamed "Anthony." Pindrop caught a real-time deepfake candidate during its own hiring process—the "Ivan X" case used throughout this report. In several cases, the operatives are being interviewed by the vendors' research teams after detection.

There are two useful reads on this pattern.

- **First, the universality of the problem:** Even firms with the best tools see it constantly in their own funnels.
- **Second, a credibility check:** When a vendor describes a fraud pattern from their own pipeline, they are speaking from primary data, not marketing.

The vendors that are best attuned to a complicated problem are often the ones who have experienced it firsthand.

6 The Candidate Fraud category materialized in roughly eighteen months—and that timing itself is the news.

A final observation: of the more than sixty vendors profiled in this landscape, at least fifteen launched or substantially rebuilt their hiring-fraud product between mid-2025 and early 2026. Many cited the same Gartner projection, the same KnowBe4 incident, the same regulatory and threat-intelligence reporting. The category did not grow—it crystallized in roughly four quarters, from a few scattered point solutions into a coherent (if fragmented) market.

This is the timing context every buyer is operating inside, whether or not they realize it. The products on the table today are, in most cases, less than two years old as fraud-specific tools—built fast, in a crisis, often by companies extending into an adjacent problem. The strongest of them are doing serious work. The weakest are riding the category's launch energy with thin substance underneath. Distinguishing the two requires the lens this report builds over the next five sections.

SECTION 1

The State of Candidate Fraud: A Fast-Evolving Threat

How Candidate Fraud Changed from a Hiring Quality Issue to an Organizational Risk

↑ CONTENTS

What Candidate Fraud Looks Like Today

Hiring teams have long managed varying degrees of misrepresentation, and for many years the consequences were relatively contained. The result was usually a poor hiring decision, additional recruiting costs, or a role that had to be reopened sooner than expected.

It was largely a quality control problem. Today's candidate fraud is far more serious.

- In July 2024, **KnowBe4**—a firm whose entire business is teaching people to spot deception—hired a North Korean operative who passed every hiring control in place before being caught by endpoint security software after his first day.
- In early 2025, **Pindrop**—itself a fraud-detection vendor—caught a candidate running a real-time deepfake face overlay during a remote interview. The same identity resurfaced eight days later through a different recruiter. The IP traced to a location near the North Korean border.
- And in 2025–2026, the security firm **Nisos** ran the experiment in reverse—deliberately advancing a suspected operative and uncovering a 40-device laptop farm in Florida running coordinated infrastructure. That single thread exposed an organized network running multiple fake identities at different U.S. companies simultaneously.

The security industry has seen this threat model before. A 2023 breach at MGM started with a ten-minute phone call to an IT help desk and ended in roughly $100 million in damage. The lesson security teams took from it: the human identity layer is the soft target.

Candidate fraud is the same attack, moved one step earlier. If an attacker can impersonate an employee to gain access, they can impersonate a candidate to get hired, and arrive inside the perimeter with a laptop, credentials, and a W-2.

Fraud in recruiting and hiring is not a new problem. The threat has been operating for years. The defensive market is months old. That asymmetry is the structural condition this report explores.

The Evidence

Evidence of candidate fraud is emerging from multiple directions. Employers are reporting suspicious hiring activity, vendors are measuring increasing levels of fraud within their own environments, and security teams are documenting more sophisticated attempts to exploit hiring processes. Gartner's projection that one in four candidate profiles globally could be fake by 2028 reflects the scale of concern surrounding the issue.

Across the market, technology providers are reporting rising levels of suspicious activity. Crosschq identified fraud indicators in 12% of applicants during 2025. Pindrop reported signs of fraud among roughly one in six applicants to its own roles. Ashby observed that approximately one in ten applicants carried multiple fraud signals, while Tofu estimates that between 20% and 30% of applications for certain remote engineering positions may be fraudulent.

1 in 4

candidate profiles globally could be fake by 2028

GARTNER

12%

of applicants showed fraud indicators in 2025

CROSSCHQ

1 in 6

applicants to Pindrop's own roles showed signs of fraud

PINDROP

20–30%

of applications for certain remote engineering roles may be fraudulent

TOFU

These figures vary because each organization measures different populations, uses different methodologies, and defines fraud differently. Despite those differences, the data points in the same direction. Organizations across industries, geographies, and hiring environments are reporting higher levels of suspicious activity and encountering fraud techniques that are more sophisticated than those seen even a few years ago.

SECTION 2

The Hiring Funnel Under Threat

Where Fraud Appears and How Organizations Are Responding

↑ CONTENTS

Candidate fraud does not occur at a single point in the hiring process. *Different threats emerge at different stages of the hiring lifecycle, creating different risks and requiring different forms of detection and prevention.*

The hiring lifecycle can be viewed as a series of interconnected control points. Early stages are vulnerable to high-volume attacks such as automated applications, synthetic candidate profiles, and AI-generated resumes. As candidates move through the process, the focus shifts toward identity verification, assessment integrity, interview authenticity, reference validation, and background screening. After hire, organizations face a different challenge altogether: ensuring the individual performing the work is the same individual who completed the hiring process.

The figure below illustrates how fraud manifests across the hiring lifecycle and the primary defenses organizations use at each stage. While candidate fraud is often discussed as a single issue, the reality is a collection of related challenges that emerge throughout the hiring process.

The sections that follow examine each stage of the hiring lifecycle in greater detail, exploring the threat landscape, the strengths and limitations of existing defenses, and the technologies emerging to address those challenges.

FIGURE 2

Fraud Across the Hiring Lifecycle

STAGE	FRAUD VECTOR	PRIMARY DEFENSE
TOP OF FUNNEL		HIGH-VOLUME THREATS
1 **Application Intake**	Automated & synthetic applications	Manual review, screening rules
2 **Resume Screening**	AI-polished & fabricated resumes	Recruiter judgment, downstream checks
MID-FUNNEL		LOWER VOLUME, HIGHER STAKES
3 **Identity Verification**	Stolen / synthetic identities, deepfakes	ID + biometric + liveness checks
4 **Assessments**	Proxy test-takers, AI assistance	Proctoring, browser monitoring *(established)*
5 **Interviews**	Deepfake overlays, proxy interviewees	Deepfake detection, interview integrity
6 **References**	Coached & coordinated references	Device fingerprinting, IP analysis
7 **Background Checks**	Records valid, person may not be	Screening *(established)* + identity verification
POST-OFFER		HIRING RISK BECOMES ORGANIZATIONAL RISK
8 **Onboarding**	The handoff	Identity reconfirmation before access expands
9 **Post-Hire**	Credential sharing, identity handoffs	Device & access monitoring *(thinnest market)*

Figure 2. Some stages benefit from established controls and growing vendor investment. Others continue to rely heavily on manual review, professional judgment, or technologies that were never designed to address today's fraud risks. Understanding where threats appear and how defenses align to those threats provides a practical framework for evaluating risk, identifying gaps, and prioritizing investments.

1 Application Intake

Candidate fraud often begins at the broadest point of entry into the hiring process. Organizations have spent years making it easier to apply, expanding candidate reach and increasing applicant volume. Those same changes have also made it easier to generate resumes at scale, automate applications, and create synthetic candidate profiles.

At this stage, the challenge is often volume rather than sophistication.

Recruiting teams may receive hundreds or thousands of applications that appear legitimate on the surface but contain fabricated information or coordinated submissions. One TA leader we interviewed described receiving approximately 1,500 applications overnight for a single engineering role and ultimately stopping review beyond a certain threshold because separating legitimate candidates from fraudulent ones became unsustainable.

Most applicant tracking systems were designed to manage applications and hiring workflows, not assess fraud risk. As a result, organizations often rely on manual review, basic screening rules, and recruiter judgment to identify suspicious activity.

> ⚠ WHAT TO KNOW
>
> The greatest risk at this stage is not simply that fraudulent candidates enter the funnel. Qualified candidates can be overlooked when recruiting teams become overwhelmed by volume and noise.

2 Resume Screening

Resume screening has always required recruiters to assess whether a candidate's experience and qualifications match the requirements of a role. Generative AI has made that task more complex by making it easier for people to create polished resumes, tailor content to job descriptions, and enhance or fabricate experience.

The challenge talent acquisition teams face is rarely determining whether AI was used. Increasingly, the question is whether the information being presented accurately reflects the candidate's skills, accomplishments, and work history. Recruiters may encounter resumes that appear highly qualified and professionally written while containing exaggerated, misleading, or entirely fabricated information.

Most resume screening processes were designed to identify candidates' baseline qualifications, not validate the authenticity of the information provided. As a result, many organizations continue to rely on recruiter judgment and downstream verification processes to identify discrepancies.

3 Identity Verification

As candidates move deeper into the hiring process, the threat begins to shift. Earlier stages are often defined by volume: large numbers of applications, resumes, and profiles that must be filtered and evaluated. Identity verification introduces a different challenge: determining whether the person participating in the hiring process is who they claim to be.

This is where organizations encounter stolen identities, synthetic identities, forged documents, and impersonation attempts. It is also where some of the most widely discussed fraud techniques emerge, including deepfake video overlays, proxy interviewees, and situations where the individual who completes one stage of the hiring process is not the same person who appears in subsequent interviews or ultimately performs the work after hire.

The volume of fraud at this stage is lower than at the top of the funnel, but the potential consequences are significantly greater. Once identity is compromised, every downstream hiring decision becomes less reliable. Organizations are increasingly adopting identity verification technologies that compare government-issued identification, biometric data, device information, and liveness checks to establish confidence that a candidate is who they claim to be.

> ⚠ WHAT TO KNOW
> Once identity is in question, every downstream hiring decision becomes less reliable.

4 Assessments

Assessments are designed to help organizations evaluate a candidate's knowledge, skills, and capabilities before making a hiring decision. They have also become a target for unauthorized assistance, proxy test-takers, AI-generated responses, and leaked question banks.

The central challenge at this stage is authenticity. Organizations need confidence that the work being evaluated accurately reflects the candidate's own abilities and not the assistance of another person or technology.

Unlike many other areas of candidate fraud, assessment integrity is supported by a relatively mature set of defenses. Proctoring, browser monitoring, plagiarism detection, and behavioral analysis existed long before the current wave of AI-enabled fraud. As a result, organizations have more established options for protecting assessments than they do for many other stages of the hiring process.

Recent data suggests those controls continue to play an important role. CodeSignal reported that **35%** of proctored assessments were flagged for suspicious activity in 2025. The company also found that non-proctored assessments produced average score inflation of **almost 8%**, compared with **less than 2%** for proctored assessments, providing a measurable example of the impact assessment controls can have on outcomes.

5 Interviews

As candidates move from assessments into interviews, the threat shifts from assistance to impersonation. Earlier stages focus on whether the application, resume, or assessment reflects the candidate's capabilities. Interviews focus on whether the person participating in the hiring process is the same individual the organization intends to hire.

This is where organizations encounter deepfake video overlays, proxy interviewees, real-time AI coaching tools, and situations where one individual completes part of the hiring process while another appears in subsequent interviews or ultimately performs the work after hire.

> *Once again, the volume of fraud at the Interview stage is lower than at the top of the funnel, but the stakes are considerably higher.*

Interviews have traditionally been one of the most trusted stages of the hiring process. Recruiters and hiring managers rely on conversation, observation, and professional judgment to evaluate a candidate's experience, communication skills, and fit for a role. But remote hiring has made many of those signals more difficult to interpret, particularly as AI-generated content and deepfake technologies become more accessible.

The challenge is that interview fraud occurs live, in a medium that most interview technologies were designed to facilitate rather than verify. As a result, organizations are increasingly supplementing interviewer judgment with identity verification, deepfake detection, behavioral analysis, and interview integrity technologies.

---

6 References

Reference checks are intended to validate a candidate's experience, performance, and professional reputation through conversations with former managers, colleagues, or other professional contacts. Their effectiveness depends heavily on trust.

That trust can be exploited. Candidates may provide references who have been coached, compensated, or otherwise prepared to confirm employment histories, job responsibilities, or performance claims. In some cases, multiple references may be connected to the same individual, organization, device, or network, creating the appearance of independent validation when none exists.

Unlike earlier stages of the hiring process, reference fraud often succeeds because the process appears legitimate. The conversation takes place. Questions are answered. Information is confirmed. The challenge is determining whether the source of that information is authentic and independent.

Several vendors have begun introducing additional validation techniques, including device fingerprinting, IP analysis, and cross-reference matching to identify potential coordination among references. These approaches can help uncover patterns that would be difficult to detect through conversation alone.

7 Background Checks

Background checks are often viewed as one of the strongest verification steps in the hiring process. They can confirm employment history, education, criminal records, professional licenses, and other information used to support hiring decisions.

The challenge is that background checks validate records, not necessarily the individual presenting them.

A candidate may successfully pass a background check because the documents, credentials, and identity being verified are legitimate. That does not guarantee the person participating in the hiring process is the individual those records describe.

The KnowBe4 workforce infiltration incident demonstrated this distinction clearly: The operative passed background screening and reference checks using a stolen identity before being detected after hire.

As identity fraud becomes more sophisticated, organizations are increasingly recognizing that background checks and identity verification serve different purposes. One confirms information associated with an identity. The other helps establish confidence that the individual presenting that identity is authentic.

Background checks remain an important control, particularly in regulated industries and sensitive roles. However, they are most effective when combined with identity verification, interview integrity measures, and ongoing monitoring throughout the hiring process.

---

8 Onboarding

For many organizations, onboarding marks the end of the hiring process. For sophisticated fraud actors, it may be the beginning.

By this stage, a candidate has successfully navigated applications, interviews, references, and background checks. They have been offered a position, granted access to systems, and issued company equipment. As a result, onboarding often provides the first opportunity to convert hiring fraud into operational, financial, or security risk.

Organizations have documented a range of onboarding-related warning signs, including laptops shipped to addresses that do not match employee records, requests to install remote-access software, unusual device activity, and inconsistencies between the individual who completed the hiring process and the person participating in onboarding activities.

This stage also introduces one of the most significant risks in candidate fraud: the handoff. The individual who completed interviews and accepted the offer may not be the same individual who ultimately receives system access or performs the work. In these cases, the hiring process may have functioned exactly as designed while still failing to prevent fraud.

Historically, onboarding has focused on equipment provisioning, paperwork, training, and access management. Organizations are increasingly recognizing that it is also an important verification point and an opportunity to confirm identity before access is expanded further.

> ⚠ **WHAT TO KNOW**
>
> Onboarding represents the transition from hiring risk to organizational risk. The closer a candidate gets to systems, data, and business operations, the greater the potential consequences if fraud succeeds.

9 Post-Hire

> *Most hiring controls are designed to answer a single question: is this person who they claim to be? But post-hire introduces a different challenge: is this still the same person?*

Historically, organizations assumed the identity verification process ended once employment began. That assumption becomes increasingly problematic in remote and distributed work environments where employees may rarely, if ever, interact with colleagues in person.

This is where organizations encounter risks such as credential sharing, unauthorized delegation of work, account misuse, and identity handoffs. In some documented workforce infiltration cases, the individual who completed the hiring process transferred access, responsibilities, or both to another person after employment began. The fraud was not the hiring decision itself. The fraud occurred after trust had already been established.

Many organizations have strong hiring processes but limited mechanisms for validating identity after onboarding. As a result, post-hire monitoring often depends on controls owned by Security and IT, including device management, access monitoring, behavioral analytics, geolocation analysis, and endpoint detection tools.

This stage remains one of the least-developed areas of the candidate fraud market, despite carrying some of the highest potential consequences. While several providers are beginning to address identity continuity and workforce verification, most organizations still treat hiring and security as separate processes with limited coordination between them.

What the Hiring Lifecycle Reveals

The hiring lifecycle reveals three important realities: Fraud emerges differently at each stage, defenses are unevenly distributed across the process, and some of the highest-impact risks occur after organizations believe verification is complete. These patterns provide a useful framework for evaluating both technology investments and operational controls.

SECTION 3

What Real Incidents Reveal

Documented Examples of Candidate Fraud

↑ CONTENTS

Candidate fraud is often discussed in terms of statistics, projections, and emerging technologies. Those data points are important, but real-world incidents provide a different kind of insight. They show how fraud occurs, where existing controls break down, and the consequences organizations face when detection fails.

No single incident represents the entire candidate fraud landscape. Some involve sophisticated nation-state actors attempting to gain access to systems and intellectual property. Others involve identity manipulation, deepfake interviews, fabricated references, or coordinated efforts to misrepresent qualifications and experience. Taken together, these cases illustrate an important reality: candidate fraud rarely succeeds because a single control fails. It succeeds when gaps emerge between controls, processes, and organizational responsibilities.

The following examples draw from public disclosures, company investigations, security research, and industry reporting. Each highlights a different point of vulnerability within the hiring lifecycle and offers lessons that extend beyond the specific incident itself.

CASE STUDY 1

KnowBe4 and Workforce Infiltration

In July 2024, security awareness company KnowBe4 disclosed that it had inadvertently hired a North Korean operative posing as a U.S.-based software engineer. The individual reportedly used a stolen American identity, an AI-enhanced photo, and successfully completed multiple interviews, a background check, and reference verification before being hired.

The incident was discovered only after the employee received a company workstation and attempted to install malware shortly after activation. Endpoint security controls detected the activity before any known damage occurred.

LESSONS FOR ORGANIZATIONS

- ✓ Traditional hiring controls can validate records without validating identity continuity.
- ✓ Fraud can succeed between hiring checkpoints, even when individual controls appear to work as designed.
- ✓ Security monitoring may detect threats that recruiting processes never identify.

CASE STUDY 2

Pindrop and Deepfake Interview Fraud

In 2025, Pindrop documented a candidate using a real-time deepfake face overlay during a remote interview. Investigators identified subtle inconsistencies in facial movement and audio synchronization before confirming the use of face-swap technology. The candidate was rejected, but the same identity reportedly resurfaced through another recruiting channel days later.

The incident highlighted how quickly fraudulent identities can move across organizations and how difficult sophisticated interview fraud can be for humans to identify through observation alone.

LESSONS FOR ORGANIZATIONS

- ✓ Visual presence is no longer sufficient proof of authenticity.
- ✓ Interviewers need clear escalation paths when something feels suspicious but cannot be immediately explained.
- ✓ Interview integrity increasingly requires a combination of awareness, process, and technology.

CASE STUDY 3

Reference Manipulation and Trust-Based Verification

Reference fraud remains one of the most common forms of candidate deception. Recruiters and screening providers continue to report cases involving fabricated references, coordinated reference networks, and individuals posing as former managers or colleagues.

Because reference checks rely heavily on trust and candidate-provided contact information, the process can appear legitimate even when the source of the information is not independent or authentic.

LESSONS FOR ORGANIZATIONS

- ✓ Reference checks validate information, but they do not always validate the source.
- ✓ Independent verification strengthens the value of reference conversations.
- ✓ References should complement other verification methods rather than serve as definitive proof.

What These Incidents Reveal

Taken together, these incidents suggest that effective candidate fraud prevention depends less on identifying a single threat and more on building a coordinated approach across the hiring lifecycle. *Organizations that understand where fraud occurs, how it evolves, and who is responsible for responding will be better positioned than those relying on isolated controls or point solutions.*

SECTION 4

Governance and Accountability

Who owns candidate fraud—and who answers for it

↑ CONTENTS

The Ownership Gap

Candidate fraud creates an unusual organizational challenge because it enters through the hiring process but rarely remains a recruiting issue. A fraudulent candidate can affect hiring outcomes, information security, compliance, business operations, intellectual property, and organizational reputation. As a result, multiple functions have a stake in the problem, yet ownership is often unclear.

Talent Acquisition is typically the first team to encounter fraud. Recruiters review applications, conduct interviews, verify information, and manage hiring workflows. They are often the first to recognize suspicious activity and the first to receive requests to evaluate new fraud prevention technologies. Yet many of the risks associated with candidate fraud extend well beyond the traditional responsibilities of recruiting.

Security teams may own identity risk, workforce infiltration concerns, and post-hire monitoring. Legal and Compliance teams may become involved when investigations, adverse actions, privacy requirements, or regulatory obligations arise. IT teams frequently support onboarding controls, device management, and access provisioning. Business leaders ultimately bear the operational consequences when fraudulent candidates enter the organization.

> *The result is an ownership gap. Candidate fraud touches multiple functions, but in many organizations no single function owns the issue end-to-end.*

This challenge appeared consistently throughout the research. Organizations of different sizes, industries, and maturity levels described remarkably similar questions:

Who should investigate suspicious activity?

Who decides whether a candidate should be removed from consideration?

Who owns the budget for fraud prevention tools?

Who is responsible for developing response procedures?

And who becomes accountable when fraud succeeds despite existing controls?

The answers were often unclear.

What the Research Revealed

One of the more surprising findings from this research was how similar the ownership challenge looked across organizations with dramatically different resources.

ONE END OF THE SPECTRUM

A team of two

At one end of the spectrum, a Talent Acquisition leader at a mid-sized security company described managing candidate fraud concerns with a team of two. He had personally evaluated emerging fraud technologies, researched building his own detection workflows using AI and automation tools, and taken responsibility for monitoring developments in the market because no formal owner existed elsewhere in the organization. The business recognized the risk, but responsibility remained largely informal.

THE OTHER END

A global enterprise

At the other end of the spectrum, a talent leader at one of the world's largest technology companies described a very different environment with remarkably similar dynamics. Security teams were aware of candidate fraud. Leadership understood the risks. Resources were available. Yet responsibility remained distributed across multiple functions, and no single team owned the issue from detection through response.

FIGURE 4

The Architectural Seam

Figure 4. Two different organizations, one structural failure. Candidate fraud sits in the same architectural seam — between Talent Acquisition and Security — at every scale.

The common challenge was not budget, organizational size, or technical sophistication. It was accountability.

If the ownership gap were simply a matter of resources, large enterprises would have solved it already. Instead, organizations at both ends of the spectrum described similar uncertainty around governance, escalation, and decision-making. The scale differed. The structure of the problem did not.

Why Security Has Historically Stayed Out

Given the security implications of workforce infiltration and identity fraud, it might seem surprising that Security is not more deeply involved in hiring processes today.

Part of the explanation is historical. Recruiting and security have traditionally operated in different domains with different objectives. Recruiting focuses on attracting talent, reducing friction, and creating positive candidate experiences. Security focuses on protecting systems, data, and organizational assets. Candidate fraud sits between those worlds.

Many security teams also view hiring as an HR responsibility and become involved only after employment begins. Likewise, Talent Acquisition teams often lack established pathways for escalating suspicious candidate activity to Security. The result is that candidate fraud can fall into a gap between functions, with each assuming the other owns part of the problem.

That dynamic appears to be changing. Vendors consistently reported increased engagement from CISOs and security organizations, particularly following high-profile workforce infiltration incidents. As organizations recognize that hiring decisions can create security risks, collaboration between Talent Acquisition and Security is becoming more common.

Governance Models Emerging Today

While no single ownership model has emerged as the industry standard, four governance approaches appeared repeatedly throughout the research.

Models Emerging Today

1 Talent Acquisition-Led Model

In many organizations, particularly smaller employers and those early in their fraud prevention journey, Talent Acquisition remains the primary owner of candidate fraud efforts. Recruiters and recruiting operations teams evaluate technologies, identify suspicious activity, document concerns, and coordinate responses. Security involvement is often limited to specific incidents or high-risk roles. This model benefits from simplicity and speed, but it can place significant responsibility on teams that may lack the resources, expertise, or authority needed to address more sophisticated threats.

2 Joint Talent Acquisition + Security Model

Many mid-market and enterprise organizations are moving toward shared ownership between Talent Acquisition and Security. Under this model, Talent Acquisition remains responsible for the hiring process and candidate experience, while Security provides expertise, investigative support, and guidance for higher-risk situations. Escalation procedures are defined in advance, allowing recruiters to focus on hiring decisions while security teams address potential fraud concerns. This approach balances hiring needs with risk management and was one of the most common models observed among organizations actively investing in candidate fraud prevention.

3 Fraud Committee Model

Larger organizations increasingly rely on formal cross-functional governance. These committees typically include representatives from Talent Acquisition, Security, HR, Legal, Compliance, IT, and in some cases Internal Audit or Privacy functions. The group reviews incidents, evaluates technologies, updates policies, monitors trends, and establishes organizational standards. While this model requires more coordination, it also creates shared accountability and reduces reliance on any single function or individual. One notable example described during the research involved a dedicated security manager embedded within the HR organization and coordinating activities across multiple departments. Rather than requiring recruiters to seek security support when needed, security participation was built directly into the operating model.

4 Executive-Sponsored Model

Smaller organizations often lack the scale necessary for formal committees or dedicated fraud teams. In these environments, executive sponsorship becomes particularly important. A senior leader explicitly assigns ownership, establishes expectations, secures funding, and ensures the issue receives appropriate visibility. The resulting program may be simpler than those found in larger enterprises, but it provides the accountability necessary to move beyond ad hoc responses.

EVERY ORGANIZATION SHOULD ASK A SIMPLE QUESTION

What would our ownership model look like *the day after a candidate fraud incident,* and what can we implement today?

The Governance Lesson

Across organizations of every size, ownership proved to be a stronger predictor of progress than technology selection. The organizations making the most progress had established accountability, escalation procedures, security involvement, and shared governance before expanding technology investments.

The next section examines how the market is responding, including the technologies, vendors, and solution categories emerging to address different forms of candidate fraud.

SECTION 05

The Market *Landscape*

Three vendor buckets, coverage by stage, and the gaps that cut across both.

↑ CONTENTS

> *The candidate fraud market looks crowded. It isn't — not in the ways that matter.*

More than sixty vendors now claim some version of fraud detection, identity verification, or hiring integrity as a core capability. What's missing is not vendors. It's coverage. The market has organized itself around the stages of the hiring lifecycle where fraud is easiest to detect, and thinned out almost entirely where it's most damaging. Understanding that pattern is more useful than counting solutions.

The sections that follow map the market three ways: by vendor structure, by hiring stage, and by the gaps that cut across both. The goal is not to identify the best vendor. It is to give buyers the lens to predict where any vendor will fall short before they sign a contract.

At the highest level, vendors fall into three buckets.

FIGURE 3 · THREE BUCKETS OF VENDORS

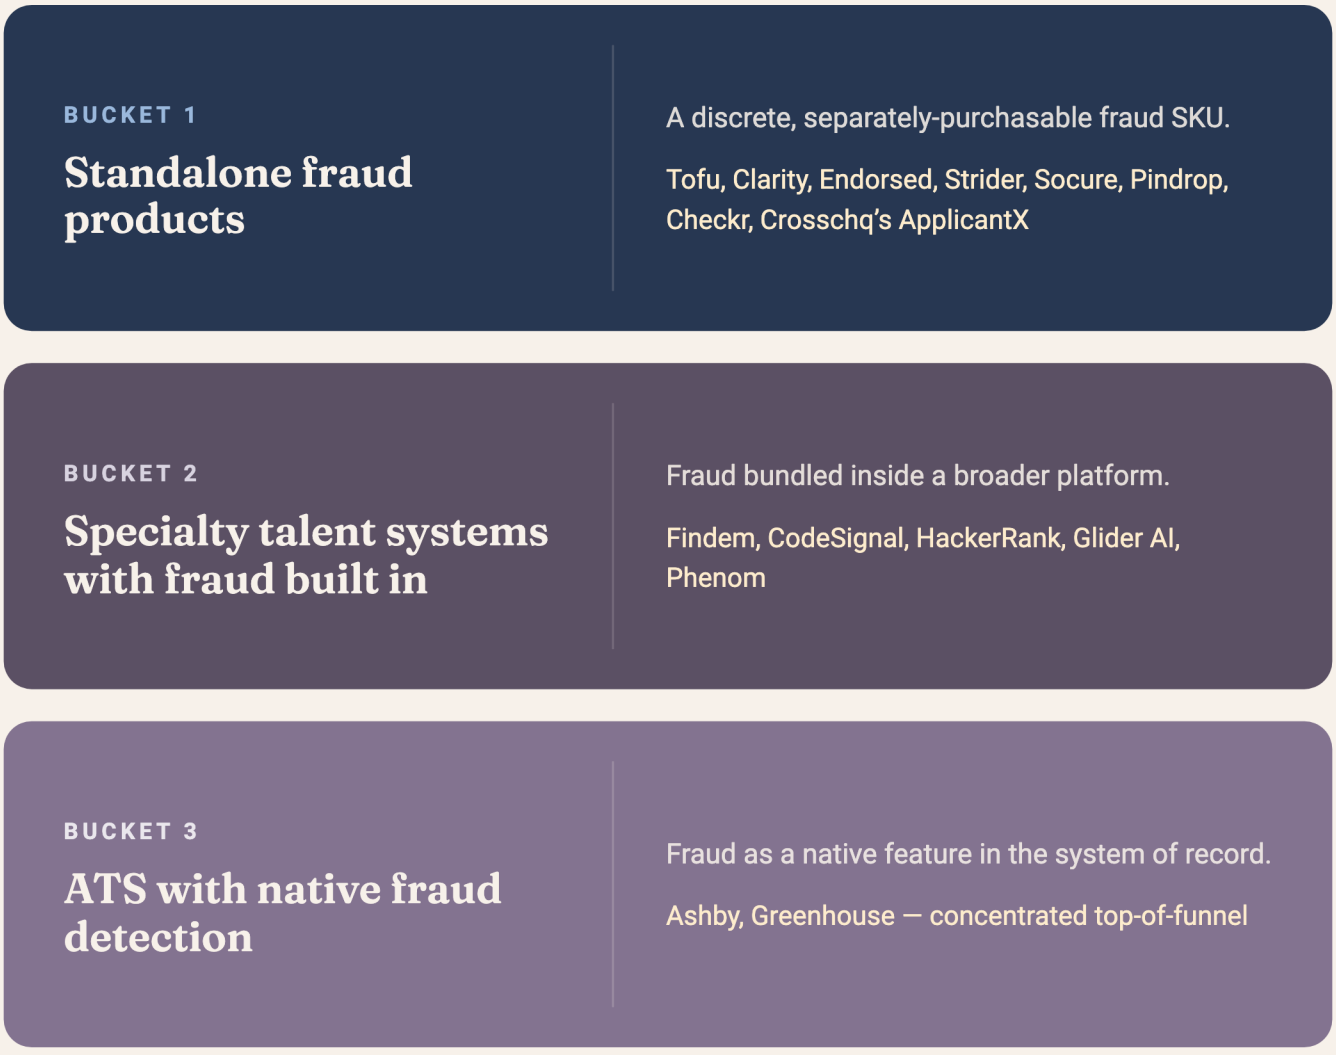

Figure 3. Three buckets of vendors offering Candidate Fraud solutions.

1 Standalone fraud products

Vendors offering a dedicated, separately-purchasable fraud product. The company may sell other things. What matters here is that the fraud capability is a discrete SKU a buyer can sign a contract for on its own. For example:

- **Tofu, Clarity, Endorsed, and Strider** sit here as purpose-built fraud systems
- **Socure, Pindrop, and Checkr** as larger identity-verification and security firms extending decade-old infrastructure into hiring processes and systems
- **Crosschq's ApplicantX**, available as either an API-first integration or a standalone enterprise app even though the parent company also sells reference checks and recruiting analytics

2 Specialty talent systems with fraud detection built in

Vendors whose fraud capability is bundled inside a broader product and cannot be purchased standalone. The buyer commits to the larger platform; the fraud capability comes with it.

- **Findem** offers the Authenticity Suite, available to all Findem customers inside the Inbound Applicants tab
- **CodeSignal, HackerRank, and Glider AI** offer cheat-detection inside the assessment platform
- **Phenom** features a Fraud Detection AI Agent inside Phenom Interview Intelligence

The fraud detection here is genuine but bounded by the host product's scope and roadmap.

3 Applicant tracking systems with native fraud detection offerings

Recruiting platforms adding fraud as a native feature, currently concentrated at the top of the funnel. Structurally this is a refinement of Bucket 2, because an ATS-native fraud capability is, by definition, bundled inside a broader product. But it earns its own bucket because the ATS is the largest single system in talent acquisition while the others are typically considered secondary systems.

The ATS is the system of record for hiring, and the fraud feature is sitting in the recruiter's daily workflow rather than a separate tool here.

At the time of this research, only two ATS providers—Ashby and Greenhouse—have shipped meaningful native fraud detection. Coverage in this bucket is focused on top-of-funnel by design, designed to surface fraud signals at the point of application rather than further down the funnel in candidate interview and assessment or new hire onboarding.

Beyond the Buckets: Solutions by Stage

The candidate fraud market is further understood through the hiring funnel. As outlined above, most solutions address a specific stage of the hiring process rather than providing end-to-end coverage.

There are several ways to evaluate vendors: Based on where fraud risk is emerging, where existing controls are weakest, and which risks would create the greatest impact if they succeeded. But when looking for the best-fit solution, it's important to note that the market is also influenced by its origins.

Identity verification providers originated in financial services and digital trust. Assessment security vendors evolved from education and testing environments. Background screening providers expanded from employment verification. Applicant tracking systems introduced fraud capabilities as an extension of recruiting workflows. Cybersecurity firms entered the market through workforce infiltration and insider-risk concerns.

As a result, organizations often encounter vendors approaching the same problem from very different starting points. Understanding where a solution originated can provide useful context for understanding the problems it solves best.

Market Coverage by Category

CATEGORY	FUNNEL STAGES	MATURITY / STRUCTURE	POTENTIAL GAPS
Purpose-Built Fraud Platforms	Application Intake → Onboarding (varies by vendor)	EMERGING Standalone	Broad coverage claims, but few vendors provide equally strong capabilities across every stage. Response workflows and compliance controls remain inconsistent.
ATS-Native Fraud Detection	Application Intake, Resume Screening	EMERGING Embedded	Limited market adoption. Most ATS platforms have not yet introduced comparable fraud detection capabilities.
Resume & AI Content Detection	Application Intake, Resume Screening	EMERGING Mixed	AI-generated content remains difficult to identify reliably. False positives and candidate experience concerns remain significant challenges.
Identity Verification (Hiring)	Identity Verification, Interview, Onboarding	EMERGING Mixed	Biometric privacy regulations and varying state requirements create implementation complexity.
Assessment Anti-Cheating	Assessment	ESTABLISHED Embedded	AI-assisted coding, interview coaching tools, and behavioral assessment manipulation continue to evolve faster than detection methods.
AI Interview Fraud Detection	Interview Screening, Live Interview	EMERGING Mixed	Limited independent validation of detection accuracy. False positives remain a concern, particularly across diverse communication styles.

Market Coverage by Category *continued*

CATEGORY	FUNNEL STAGES	MATURITY / STRUCTURE	POTENTIAL GAPS
Video Interview Platforms	Interview Screening, Live Interview	ESTABLISHED Embedded	Most platforms provide passive evidence trails rather than active detection in real time.
Deepfake Detection	Live Interview, Identity Verification	EMERGING Standalone	Real-time video detection remains technically challenging. Accuracy varies based on media quality and attack sophistication.
Reference Verification	Reference Check	ESTABLISHED Mixed	Traditional reference processes remain vulnerable to fabricated references, proxy references, and coordinated fraud networks.
Background Screening & Onboarding Verification	Background Check, Onboarding	ESTABLISHED Embedded	Most providers focus on historical verification and have limited visibility into interview fraud, identity continuity, or post-hire handoffs.

Several themes emerged consistently throughout this research.

Buy-Side Innovators

Some talent acquisition organizations are choosing to build their own fraud solutions rather than buy. In one example, a Talent Acquisition leader who was dissatisfied with the products currently available began developing a detection workflow using Claude Code, Zapier, and commercially available AI tools. Similar examples have started to emerge across the market.

While these efforts are unlikely to replace dedicated fraud prevention platforms long term, they highlight both the accessibility of modern AI tools and the immaturity of the current market. When practitioners conclude that assembling their own solutions is preferable to purchasing available products, it provides insight into the gaps that still exist between organizational needs and vendor capabilities.

This approach is best viewed as a temporary bridge rather than a permanent strategy, but it underscores how rapidly both buyers and vendors are experimenting in this space.

Human Judgment Remains Central

Despite approaching the problem from different directions, most vendors position fraud detection as a decision-support capability rather than an automated hiring decision. Vendors consistently described their role as surfacing signals and evidence while leaving hiring decisions to people.

This human-centric approach reflects both practical and legal realities. Fraud detection remains imperfect, false positives create business and compliance risks, and hiring decisions ultimately require human judgment.

The Solution Density Gradient — and What It Tells Us

Lay the available tooling over that lifecycle and the imbalance is stark:

Application intake and identity verification are well-served (15+ vendors identified)

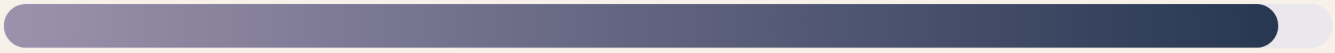

Real-time deepfake detection at the interview is served by a smaller cluster

Reference checking is essentially a two-vendor field

And post-hire continuous verification is the thinnest segment in the entire market

> *This is the funnel's central finding: Defenses cluster where fraud is easiest to detect and cheapest to act on, and thin out exactly where the most damaging fraud actually succeeds.*

The top of the funnel is crowded because top-of-funnel signals are inexpensive to gather and low-friction for the candidate. Post-hire is empty because post-hire monitoring is operationally hard, organizationally ambiguous, and nobody's obvious job.

Two structural patterns run quietly through the rest of this report, rather than as sections of their own.

THE FIRST — A PROBLEM OF TIMING

Identity verification—arguably the highest-assurance check available—typically runs at the background-check stage, after two to four rounds of interviews with a candidate who may not be real. **The check is not failing. It is firing late.**

Both vendors focused on this stage described the same instinct in their own product design: Keep friction light at the top and hold the heavier identity ask until later in the funnel, when the cost of a bad outcome is higher and the request is more explainable to the candidate.

That logic is sound for candidate experience. It also means the most reliable check happens after the most expensive investment has already been made.

THE SECOND — A PROBLEM OF ECOSYSTEM ARCHITECTURE

No single product covers the entirety of this lifecycle. The market has fragmented by stage, and the vendor who solves application intake is rarely the vendor who solves the interview, who is rarely the vendor who solves post-hire. Which means every real-world defense is a stack—assembled from multiple vendors, with seams between them—and the question of who owns those seams is one we will return to repeatedly.

For now, the takeaway is simpler: *Fraud is a lifecycle problem. The tooling is not yet a lifecycle solution.* And the space between those two facts—concentrated at the back half of the funnel—is where the rest of this analysis lives.

CLOSING

Conclusion

The hiring funnel is now an attack surface. The decision that matters is who answers for it.

↑ CONTENTS

A research report like this one is, in part, a long answer to a short question: what should a hiring organization actually do about candidate fraud?

If the only thing a reader carries away from these pages is a list of vendors to evaluate, the report has failed at its own job. The vendors matter, and the funnel matters, and the architectures matter—but the hardest truth this research surfaced is not about any of those.

> *The hardest truth is that the unsolved problem here is organizational, not technical.*

Every vendor interviewed for this report, from opposite ends of the architecture spectrum, arrived at that conclusion independently. Every practitioner interviewed—from a two-person TA team to a sophisticated enterprise—described the same gap from a different vantage point. The technology is being built quickly enough. The threat is being studied carefully enough. The market has more capability than the average organization is using. What is missing is the organizational scaffolding that would let the capability actually do its work—the owner, the protocol, the cross-functional structure, the budget that comes from the right place, the recruiter who is enabled rather than blamed.

That gap is not a small one to close. But it is also not, in the end, a procurement problem.

Three numbers, more than any others, recurred across the research, and together they describe the shape of the problem better than any chart could.

50%

Roughly the rate at which humans correctly identify AI-generated content and real-time deepfakes, per multiple academic studies (Groh et al., PNAS 2022; Köbis et al., iScience 2021)— meaning a recruiter trying to spot a sophisticated fraudulent candidate by sight is, statistically, flipping a coin. Training recruiters to be the detection mechanism is no longer a viable strategy. It was already weak. It is becoming non-functional.

Four

The number of video interviews a North Korean operative successfully passed at KnowBe4— a company whose entire business is teaching people to spot deception. If the firm that trains the world in detecting human-vector attacks could not, internally, detect this one, the implication is not that KnowBe4 failed. The implication is that the threat exceeded what unaided human observation can do, even when the observers are professionals at exactly this kind of observation.

Zero

The number of vendor-neutral, publicly available, broadly adopted candidate-fraud response playbooks. There is a great deal of vendor-specific guidance, a few well-documented case studies, and a slowly growing body of conference talks. There is no widely accepted protocol that a TA leader at a mid-sized company can pick up tomorrow and run. The market has produced detection without producing response.

Read together, those three numbers say something specific. **Recruiters cannot reliably see sophisticated fraud—not at scale, not anymore**. The organizations being targeted include the ones best-prepared to defend themselves. And the institutional knowledge of how to actually respond, once fraud is detected, has not yet been written down.

Recruiters carry the floor. *They cannot carry the ceiling.*

That sentence is the closest this report comes to a recommendation, and it is offered as a reframing rather than a prescription. For most of the past two years, the implicit model of fraud defense has been: train the recruiter to spot it, give the recruiter a tool to flag it, hold the recruiter accountable when something slips through. That model has produced motivated recruiters operating in bubbles—the pattern this research found at every company size, from the smallest TA team to one of the largest security firms in the world.

The model that survives the next two years is more honest about the spectrum. A trained recruiter can carry the floor—the "little f" fraud that has always existed and that thoughtful human attention is still good at catching. A trained recruiter cannot carry the ceiling—a nation-state operative running a real-time deepfake, the threat profile that has outpaced unaided human observation.

The recruiter's structural job, especially in the gray zone between those two, is to escalate into a structure built to handle it: a defined cross-functional group, a written protocol, budget owned at the right level, security partnership that arrived before the incident rather than after. That structure does not exist in most organizations today. Building it is the work.

The technology will keep getting better. The vendors will keep consolidating, the embedded layer will keep commoditizing, the post-hire category will keep forming, the compliance reckoning will keep coming. None of that, by itself, closes the gap this report has been circling. The question that actually determines outcomes—and the one this entire research project kept returning to—is who owns the seams between the layers, and who adjudicates a flag once it fires.

> *That question has an answer in every organization. The answer is almost never written down.*

The hiring funnel is now an attack surface. The defenses are coming online, unevenly, faster in some places than others. The market is being built in real time, with all the unevenness that implies. But the single most consequential decision a hiring organization will make about candidate fraud in the next two years is not which vendor to buy. It is who, on the org chart, is responsible when the next one gets through.

A RESEARCH REPORT BY KYLE & CO.

About Kyle & Co

Kyle & Co is a modern research and advisory firm helping HR and talent technology leaders make smarter decisions, faster.

By blending practitioner insight, rigorous research, and strategic advisory, we uncover the signals that matter most—then translate them into action. Our work spans custom research, workforce strategy, and market analysis, supporting solution providers and HR leaders alike. Kyle & Co is the team behind the Human-Centric AI Council, the Transformation Realness podcast, and industry-shaping reports on quality of hire, workforce planning, and responsible AI.

We work with clients like Workday, ADP, GoodTime, Findem, and SmartRecruiters.

Learn more at www.kyleandco.com

CONTACT US

www.kyleandco.com
info@kyleandco.com
linkedin.com/company/kyle-co