



White Paper

Candidate Fraud in Enterprise Hiring

A Risk Management Framework
for Contingent Workforce Leaders

March 2026

Table of Contents

When Did Hiring Start Requiring Detective Work?	3
Executive Summary	5
Section 1: The New Fraud Landscape	6
Three Threat Tiers	6
The Fraud Taxonomy: Six Mechanisms in Talent Acquisition	7
Section 2: The Real Cost Of Fraud	13
The Fraud Cost Continuum	13
Section 3: The Hiring Lifecycle Defense Model	17
Integration Across Phases – And Across Partners	22
Section 4: Assessing Your Security Posture	23
1. What Is Your Attack Surface?	24
2. What Does Your Incident History Tell You?	25
3. Where are the Vulnerabilities?	26
4. Can You Close the Gaps Internally - or Do You Need to Source Capabilities?	27
Section 5: Fraud Posture Self-Assessment	28
Scoring Your Posture	30
Closing	31



When Did Hiring Start Requiring Detective Work?

Candidate fraud is not a new problem. Resume embellishment, inflated credentials, and exaggerated experience have been part of hiring for decades. What is changing fundamentally is the infrastructure behind the fraud – organized networks running coordinated placement operations, AI tools that generate polished resumes and coach candidates through live interviews in real time, and a remote work environment that removed many of the in-person signals hiring managers used to rely on without replacing them with anything equivalent.

That's why candidate fraud persists – and why it is getting harder to catch. Sometimes it is blatant. More often, it is ordinary. A resume maps almost perfectly to the job description. Employment history is technically verifiable, but not the history you interviewed against. A worker's location or availability patterns do not match what was agreed. In many cases, the person can do the work. The misrepresentation is about who they worked for, where they are working from, or how the work is getting done.

Most organizations are not built to detect this. Hiring is built on a foundation of good faith. A manager reviewing a candidate wants to believe the person in front of them is who they say they are. When someone presents well, communicates confidently, and has a resume that fits the role, the instinct is to lean in, not to interrogate. That instinct is not naivety – it is how functional working relationships begin. But it is also the gap that candidate fraud is specifically designed to exploit. Hiring managers engage in hiring candidates occasionally, not full-time, and they see only a small sample of resumes in any given year. When something looks great on paper, pressure is real, small inconsistencies get rationalized, and they only make sense in hindsight, after the cost has already been absorbed.

Sophisticated operations and high-impact scenarios do exist, especially in remote environments. But the most common exposure does not require a mastermind threat. It comes from simple gaps in verification, handoffs between teams, and hiring processes designed for speed rather than detection.

Contingent workforce programs sit at the center of this exposure. Contractor hiring channels are generally more open than core employee recruitment. Placement velocity is higher, verification windows are shorter, and multiple intermediaries create handoff gaps that sophisticated actors exploit.

Employer of Record (EOR) and Agent of Record (AOR) programs carry particular risk: client hiring managers identify and select their desired workers but typically lack the fraud detection capabilities present in a staffing vendor's recruitment organization or the client's own permanent hiring function. The result is a higher-risk surface area that most organizations have not assessed against the current threat environment.

There are a few reasons for this increased risk:

1. Managers are often under pressure to fill roles to hit outcomes. Candidate resumes that are targeted for a role will look promising to hit these outcomes.
2. Managers only see resumes for their current roles, not the broader hiring market, and therefore they don't see the common trends that staffing agencies would see over thousands of candidates.
3. Trust: People will generally trust that what they are hearing is true.
4. Their own bias can come into play; if a person was referred by a colleague, or they have strong credentials, etc.

This paper is written for leaders accountable for contingent workforce programs, whether that responsibility sits in HR, Talent, Procurement, or a dedicated workforce function. It provides a framework for understanding the six primary fraud mechanisms now active in recruitment, a model for assessing true cost exposure, a practical defense architecture, and a self-assessment tool to evaluate your current posture.

The organizations best positioned to manage this risk are those that recognize candidate fraud as an enterprise risk management challenge – not a hiring inconvenience – and invest accordingly.



About the **Author**

Kent McCrea, **CEO & President at Procom**

Kent McCrea has over 15 years of experience leading one of Canada's top staffing firms, delivering consulting and workforce solutions to many of North America's largest organizations. As AI capabilities become more advanced and widely adopted, he brings a unique perspective on the evolving staffing landscape, with deep insights into emerging market trends and industry shifts.

With contributions from
Wendy Kennah,
Chief Operating Officer at Procom

Executive Summary

The threat is systematic. Candidate fraud now operates across three tiers – opportunistic individuals, organized commercial operations, and state-sponsored actors – using six distinct mechanisms that exploit gaps across the hiring lifecycle.

Contingent programs are disproportionately exposed. Higher placement velocity, shorter verification windows, multi-party handoffs, and limited fraud detection at the hiring-manager level make contractor channels a preferred target for organized fraud groups. EOR/AOR programs, where client managers select workers without access to dedicated screening infrastructure, carry elevated risk.

The cost is larger than most organizations recognize, because the most common losses are the least visible. Routine hiring friction is absorbed without measurement. Payroll fraud losses are categorized as bad hires. Operational harm from overemployment or unauthorized subcontracting is attributed to performance

problems. Only catastrophic incidents – IP theft, regulatory penalties, ransomware – force organizations to confront the fraud behind them. By then, a single placement has triggered millions in downstream damages.

Defense requires an integrated approach across the hiring lifecycle. Effective fraud mitigation layers controls based on technology, process, and human expertise across the four phases of the hiring lifecycle – pre-hire, during-hire, on assignment, and post-hire. Organizations that concentrate verification in a single phase or rely on a single controls modality create the predictable gaps that sophisticated fraud operators will seek to exploit.

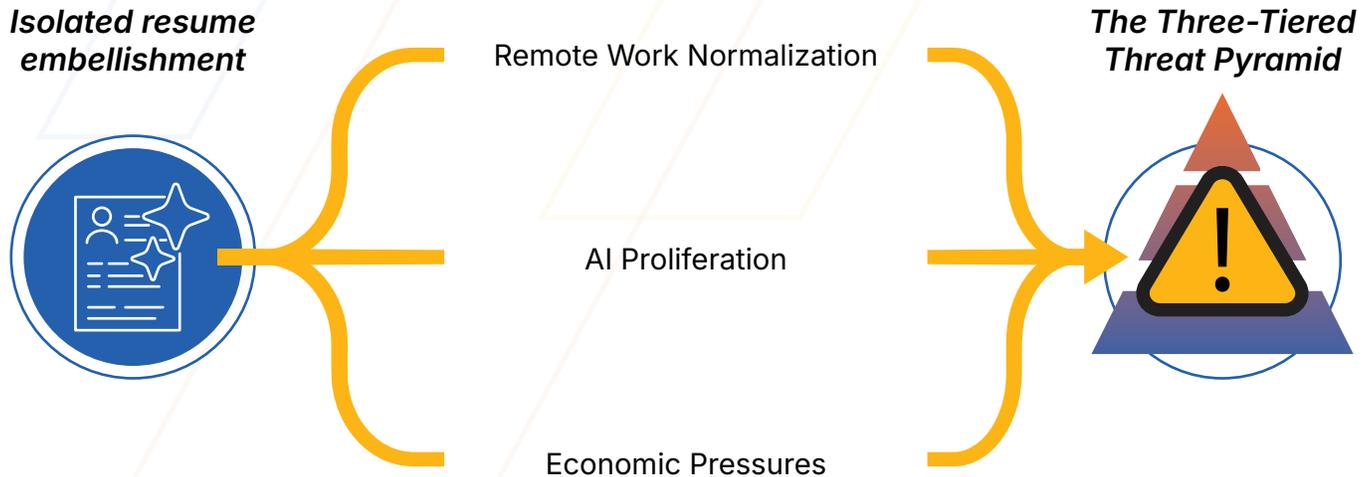
Start with your risk posture. The self-assessment framework in Section 5 enables organizations to evaluate their current controls against the threat environment and identify where gaps create the highest exposure.

Key Frameworks in This Paper

This paper introduces five interconnected frameworks:

- 1. Three Threat Tiers** – classifying fraud actors by sophistication and impact;
- 2. Six Fraud Mechanisms** – a taxonomy of attack patterns active in recruitment;
- 3. The Cost Continuum** – four levels of organizational exposure from friction costs to catastrophic loss;
- 4. The Hiring Lifecycle Defense Model** – a 4-phase × 3-modality architecture for integrated fraud controls;
- 5. The Fraud Risk Posture Self-Assessment** – a diagnostic tool for evaluating current defenses.

Section 1: The New **Fraud Landscape**



Key Insight:

Most organizations design hiring processes to catch opportunistic individuals (Tier 1). These controls provide minimal defense against organized fraud operations (Tier 2) and virtually no protection against state-sponsored actors (Tier 3). Contingent workforce channels – with their higher volume, shorter verification cycles, and multi-party handoff gaps – are where this mismatch is most exploitable.

Three Threat Tiers

Tier	Actor Type	Motivation	Sophistication	Detection Difficulty	Business Impact
Tier 1	Opportunistic Individuals	Personal financial gain, career advancement	Low to Medium	Relatively Easy	Low to Medium
Tier 2	Organized Operations	Commercial fraud at scale	Medium to High	Moderate	Medium to High
Tier 3	State-Sponsored Actors	Weapons funding, espionage, supply chain attacks	Very High	Very Difficult	Severe

The Fraud Taxonomy: Six Mechanisms in Talent Acquisition

These are not new fraud types. They are established attack patterns – perfected over decades in financial services, supply chains, cybersecurity, and government contracting – now migrating into recruitment. Each mechanism below includes the industry where

it was pioneered and where the most mature countermeasures exist. These parallels are practical: the lessons from those domains directly inform what works in defending against the recruiting variants.

1 Candidate Credential Fabrication

Falsification of credentials, history, or qualifications.

2 Identity Fraud / Impersonation

Person interviewed is not the person who performs the work.

3 Unauthorized Subcontracting / Worker Substitution

Worker secretly delegates some or all work to undisclosed third parties.

4 Overemployment / Undisclosed Conflicts

Workers simultaneously hold multiple full-time positions without disclosure.

5 Location / Access Misrepresentation

Workers misrepresent their physical location.

6 Capability Inflation / AI-Assisted Misrepresentation

Candidates misrepresent their actual skill level through AI-assisted assessments.

1. Candidate Credential Fabrication

Industry Parallel: Counterfeit goods fraud (industrial supply chain)

Definition: Falsification of education credentials, professional certifications, employment history, or technical qualifications.

Attack Surface: Pre-hire phase (resume screening, application review, initial qualification)

These cases of misrepresentation are typically conducted by a single actor and may not be malicious; however, they are still a form of fraud.

This can look like a candidate that has the skills on paper, but potentially not the right industry experience, or they have overseas experience that may not carry over to North America.

In other cases, the resume, references, sample projects, credentials, and education are fabricated, or worse, stolen. The person they say they reported may be a friend, and workplaces reported in the employment verification system do not line up the resume.

Detection Indicators:

- Certification numbers that validate in format but aren't in official issuer databases
- Employment at companies with minimal digital footprint or brief existence
- Vague or inconsistent technical explanations during deep-dive discussions
- LinkedIn connections primarily to others with similar employment patterns
- References who cannot provide specific examples of work performed

2. Identity Fraud / Impersonation

Industry Parallel: Identity theft and account takeover fraud (financial services)

Definition: The person interviewed is not the person who performs the work – this can be caused by interview proxies, deepfake-enabled video substitution, post-hire bait-and-switch, and synthetic personas built from stolen personal information.

Attack Surface: Interview phase through initial work period

Identity fraud exploits the gap between virtual interviews and actual work. Interview proxies are the simplest variant. Real-time deepfake technology now enables face-swapping during video calls with synchronized lip movement, defeating both visual verification and camera-on policies

Scenario – Wage Theft: A firm hires a cloud infrastructure engineer with strong credentials. Within weeks, the new hire struggles with basic tasks and avoids video calls. After six weeks, termination for performance. Post-exit investigation reveals a stolen identity – the person who worked the role was not the person who passed verification. Total cost – wages under a fraudulent identity, recruiting, onboarding, project delays, re-hiring – exceeds \$65,000 with no practical recourse.

Scenario – Espionage and Regulatory

Exposure: A SaaS company hires a senior developer after three strong video interviews. Post-onboarding red flags accumulate: weaker written communication, persistent audio-only calls, work patterns suggesting a different time zone. Six months later, a security audit discovers code repositories systematically cloned to external servers. The hired identity was stolen; the actual worker operated overseas. Total damages – system re-architecture, regulatory penalties, client notification – exceed \$8 million.

Detection Indicators:

- Persistent avoidance of video calls or in-person meetings post-hire
- Unable to use company-issued equipment
- Communication quality inconsistent with interview performance
- Work patterns inconsistent with stated time zone
- Technical knowledge gaps that emerge after hire despite strong interview
- Unusual VPN usage or access from unexpected geographic locations or roaming IP addresses

Real-World Context

In 2024, the U.S. Department of Justice charged multiple individuals in connection with schemes to place foreign IT workers in remote positions at U.S. companies using stolen American identities. Prosecutors described operations generating millions of dollars in fraudulent wages, with proceeds in some cases funneled to support sanctioned foreign programs. These cases represent the convergence of identity fraud, location misrepresentation, and state-affiliated objectives described across Mechanisms 2, 5, and the Tier 3 threat model.

3. Unauthorized Subcontracting / Worker Substitution

Industry Parallel: Shell company fraud (financial services); subcontractor substitution (government contracting)

Definition: A hired worker secretly delegates some or all work to undisclosed third parties. The client believes they're receiving work from a known and vetted individual, but output comes from unknown, unvetted parties operating under no contractual obligation.

Attack Surface: Post-hire (assignment execution phase)

Subcontracting operates on several motivations: capability arbitrage (the hired worker can't deliver but acts as middleman), cost arbitrage (a domestic developer subcontracts offshore or onshore, pocketing the rate difference), and access harvesting (providing undisclosed third parties access to client systems). In every variant, unknown individuals with no background checks and no contractual obligations access production systems through shared credentials – invisible to the client.

Detection Indicators:

- Simultaneous system access from multiple geographic locations
- Inconsistent coding style or approach across work products
- Work submitted during unusual hours for claimed time zone
- IP addresses associated with VPN services or data centers
- Communication style or language proficiency inconsistent with interview performance
- Refuses to turn on camera in meetings
- Won't pick up calls and insists on calling back to give time to have right person on the call

4. Overemployment / Undisclosed Conflicts

Industry Parallel: Insider threats (cybersecurity); conflict of interest fraud (corporate governance)

Definition: Workers simultaneously hold multiple full-time positions without disclosure, creating divided attention, schedule conflicts, and potential conflicts of interest – distinct from authorized moonlighting or disclosed part-time work.

Attack Surface: Post-hire (ongoing assignment management)

The risk exists on a spectrum. At the lower end, undisclosed secondary employment creates time and attention conflicts that degrade performance. At the higher end, simultaneous

roles at direct competitors create information leakage and potential breach of fiduciary duty. Even without malicious intent, overemployed workers facing simultaneous critical incidents at multiple employers cannot fully respond to either, and the resulting shortcuts – superficial code reviews, documentation gaps, security workarounds – compound over time.

AI tools have added complexity to this picture. Proficient AI users can complete work in a fraction of the time expected under pre-AI productivity baselines, making it possible to satisfy the output demands of multiple roles within a standard working week. This creates a genuinely unresolved question: where a worker delivers full-quality output across two roles, the harm may lie primarily in the



non-disclosure rather than the performance. The more immediate operational risk is the inverse – managers who are not yet calibrated to what AI-assisted workers can produce may be setting timelines that leave significant capacity undetected and unfilled for the organization paying for it. Whether AI-enabled overemployment degrades output or simply redistributes surplus capacity, the non-disclosure remains the threshold issue.

The scale of the problem is documented. In a 2021 article, the Wall Street Journal profiled the overemployed phenomenon, interviewing white-collar workers secretly holding multiple full-time remote jobs and verifying their claims through offer letters, concurrent pay stubs, and corporate emails. Workers reported earning \$200,000 to \$600,000 annually across simultaneous positions. Since then, the practice has scaled: dedicated Reddit communities now exceed 500,000 members, where participants openly coach each other on managing conflicting meeting calendars, defeating remote work monitoring software, evading employment verification databases, and avoiding detection by employers.

Detection Indicators:

- Consistent unavailability during certain hours without explanation
- Frequent last-minute meeting cancellations; persistent requests for audio-only participation
- Unable to join meetings or phone calls on short notice
- Work quality declining over time with deliverables that seem rushed
- LinkedIn or professional directories listing different current employer than records show
- LinkedIn presence absent entirely, or employment history shifting without explanation
- Unusual working hour patterns – rapid-fire activity followed by long inactivity

Employer Response

Multiple large employers have conducted bulk terminations for undisclosed moonlighting and use of activity-simulation software – incidents significant enough to draw coverage from major business outlets. Equifax developed a commercial product (Talent Report Work Inform) specifically to detect overemployed workers, signaling that the problem has reached sufficient scale to support a dedicated detection market.

5. Location / Access Misrepresentation

Industry Parallel: Supply chain provenance fraud; export control violations

Definition: Workers misrepresent their physical location, performing work from undisclosed or unapproved locations – creating data sovereignty violations, export control exposure, and information security risks.

Attack Surface: Interview, onboarding, and ongoing operations

Location determines employer payroll tax obligations, which privacy laws apply (GDPR, HIPAA), whether export control regulations are violated (ITAR, EAR), and whether client contractual requirements are met. A worker accessing EU citizen data from a non-approved jurisdiction violates GDPR regardless of the employer's knowledge. A worker performing ITAR-controlled defense work from outside the United States creates criminal liability.

The infrastructure enabling location deception has scaled significantly. Laptop farms – physical facilities housing dozens of computers operated remotely – enable workers to appear locally based while operating from overseas. Residential proxy networks route internet traffic through genuine residential IP addresses in target locations. Detection requires active monitoring, not passive assumption.

Detection Indicators:

- IP address geolocation inconsistent with stated work location or changes frequently
- Login times consistently misaligned with claimed time zone
- Video background inconsistencies (lighting, architecture, seasonal cues)
- Device shipping address different from stated location on resume
- Network latency patterns inconsistent with claimed proximity



6. Capability Inflation / Ai-Assisted Misrepresentation

Industry Parallel: Academic fraud (education sector); qualification fraud (professional licensing)

Definition: Candidates misrepresent their actual skill level through AI-assisted assessments, coached interviews, and take-home tests completed by others – passing screening for roles they cannot independently perform.

Attack Surface: Assessment and interview phase

AI tools have shifted what candidates can claim and demonstrate during a hiring process. Take-home technical assessments and asynchronous interviews, once reliable signals of capability, can now be completed with significant AI assistance and make it harder to distinguish genuine expertise from a well-prompted result.

In practice, this rarely resembles outright fraud. More often, it reflects candidates whose confidence in AI-augmented output has outpaced their independent ability. The gap becomes visible not through surveillance, but through the natural progression of the hiring process itself.

Detection Indicators:

- Inability to explain or modify their own prior work in real-time
- Eye movement or audio artifacts during video interviews suggesting off-screen prompting
- Performance degradation in proctored environments, network-restricted settings, or in-person collaboration
- Work product quality inconsistent across assignments – a signature of multiple production sources



Section 2: The Real Cost Of Fraud

Beyond The “Bad Hire” Calculation

Candidate fraud doesn't announce itself. It surfaces as a bad hire, a delayed project, a budget overrun, or an unexplained data incident – categorized and closed without the underlying cause ever being identified. The result is that most organizations are absorbing real financial losses from fraud while measuring them as something else entirely.

The costs fall into four bands, each with a distinct probability and a distinct response requirement.

The Fraud Cost Continuum



Level 1: Friction Costs

Common – Most organizations absorb these regularly

A fraudulent candidate who clears screening, onboards, underperforms, and exits looks identical to a bad hire. The candidate and recruiting agency gets paid. The replacement search begins. No fraud report is filed because no fraud was recognized.

The direct cost of a single failed placement – recruiting fees, onboarding time, project disruption, and re-fill – typically runs **\$20,000–\$50,000**. In programs placing 50 or more workers annually, this happens multiple times per quarter. It doesn't feel like fraud exposure. It feels like hiring variability. That misclassification is the problem: it means the root cause is never addressed and the pattern repeats.



Level 2: Direct Financial Loss

Periodic – Organizations making 50+ placements annually

A candidate with fabricated credentials or a falsified identity collects three to four pay periods before inconsistencies surface – typically \$15,000–\$35,000 in direct payroll theft per incident. Add the recruiting fee already paid, onboarding costs absorbed, and the cost of the replacement search, and a single incident runs **\$40,000–\$75,000** in recoverable loss.

Fraud operations target this band deliberately. The amounts are large enough to be profitable at volume but modest enough that most procurement teams close the incident rather than investigate the pattern. And that is exactly what organized fraud rings count on: each incident is handled individually – flagged as a bad hire, processed as an early termination, escalated briefly, then closed. The next incident, six weeks later with a different vendor and a different role, gets the same treatment. No one connects them because no one is looking across incidents rather than within them.

The tell is almost never a single event. It is the cluster – multiple sudden exits within a quarter, a pattern of candidates who pass screening but fail within the first sixty days, unexplained performance failures concentrated in a particular skill category or sourcing channel. Organizations that track these incidents in isolation will keep absorbing the losses. Those that track them as a program-level pattern will find the source.

Not all Level 2 fraud surfaces through failed placements, however. A candidate who fabricated their employment history but genuinely possesses the skills to perform the role may never be detected at all. Their identity is real and verifiable. Their references check out – often because they are personally known to the people listed, or because the overseas roles they actually held are difficult to independently verify. What is false is the specific experience the hiring decision was based on: the ten years of local enterprise experience that justified the rate, the domain-specific roles that made them the preferred candidate over others, the career progression that suggested institutional depth they may not have.

They onboard, perform adequately, and remain placed. No exit, no complaint, no incident. But the organization made a hiring and compensation decision on a foundation that was never accurate – potentially paying a senior rate for mid-level experience, bypassing candidates with legitimate local credentials, or placing someone in a trust or access tier their actual background would not have supported. The fraud does not announce itself because nothing visibly fails. It simply sits inside the program, undetected, for as long as the placement continues.





Level 3: Operational and Competitive Harm

Less frequent – But significantly more costly per incident

An overemployed worker billing full-time at a senior rate while splitting attention across two or more roles creates multiple compounding costs, and most organizations only see the first.

First is paid capacity that never fully shows up. The work may be good, but it is not full-time engagement. What is missing is the availability for urgent requests, sustained focus on complex problems, and the day-to-day presence a senior placement is meant to provide. As attention is divided, context-switching across similar but different environments increases the likelihood of subtle errors. These are not “incompetence” mistakes. They are the mistakes of a capable person operating at the edge of cognitive bandwidth.

Second is opportunity cost - this it is the cost most programs never put a number on. A senior placement is not priced for deliverables alone. It is priced for the impact that a fully-engaged, highly-capable person in that seat should drive over the course of a year. When that capacity is silently redirected to another employer, the organization is not receiving a discount, it is funding someone else’s second income while carrying the full cost of the role.

The third cost is the one that appears nowhere in most fraud frameworks: lost knowledge transfer. A senior worker operating at divided capacity is not mentoring junior developers, participating in architecture discussions, or bringing accumulated expertise into the day-to-day work of the team around them. In an environment where AI is accelerating the pace of delivery, the gap

between a senior person who is genuinely present – guiding how AI tools are used, reviewing what they produce, elevating the capability of the people around them – and one who is physically absent while appearing to bill full-time is not just a capacity gap. It is a compounding capability gap that affects every person on the team who should have been learning from them. Even if the work quality is exceptional, but the candidate is only working 20 hours instead of 40, it begs the question – what could they achieve with the full-time hours?

AI makes all three costs harder to detect and harder to attribute. A proficient AI user can produce convincing output in a fraction of the billed hours, which means the quality signal that would normally indicate reduced engagement may not appear at all. But there is a harder question underneath the detection problem: if a senior worker is delivering acceptable output using AI in twenty hours a week, the organization is not getting what it paid for – it is getting a fraction of what that person could produce if fully committed. The output looks adequate against expectations set before AI changed what was possible.

The fully-loaded cost of a single senior-level overemployment incident commonly runs **\$150,000–\$300,000** in salary waste, rework, and lost output. In programs with multiple concurrent placements at this level, it becomes a material budget exposure. In some cases, it also raises the risk of unknown third parties accessing client systems and data, triggering investigation costs and potential penalties.





Level 4: Catastrophic Loss

Rare – But existential when it occurs

IP theft, ransomware facilitated by a planted insider, export control or data sovereignty violations from location misrepresentation, or state-sponsored infiltration. Damages run into millions or tens of millions: regulatory penalties, legal settlements, client contract terminations, and reputational harm that takes years to rebuild. The asymmetry is stark – **a single \$50,000 fraudulent placement can trigger \$5+ million in downstream damages.**

Here is another option: IP theft, ransomware facilitated by a planted insider, export control or data sovereignty violations from location misrepresentation, or state-sponsored infiltration through a placed worker. Damages run into millions or tens of millions: regulatory penalties, legal settlements, client contract terminations, and reputational harm that takes years to rebuild.

What makes Level 4 categorically different from the levels above is not just the scale of the damage – it's the uncertainty. In every other scenario, the organization knows what went wrong. A bad placement exits. An overemployed worker is identified. A falsified resume is uncovered. The incident has edges. It can be scoped, investigated, and closed.

At Level 4, that clarity is often absent. When a worker with undisclosed affiliations or a fabricated identity has had access to systems, codebases, client data, or internal communications for months or years, the organization cannot fully determine what was touched, what was copied, what was observed, or what was communicated externally. The investigation itself becomes a significant cost – forensic analysis, legal review, regulatory notification obligations, and client disclosure – before a single dollar of direct damage is counted. And because the full scope of access may never be reconstructed with confidence, the organization is left managing residual risk it cannot precisely quantify.

This ambiguity is what drives the asymmetry. A single \$50,000 fraudulent placement can trigger \$5 million or more in downstream damages – not because the fraud was sophisticated, but because the access it enabled was broad, the exposure period was long, and the organization has no reliable way to bound what was compromised. The response is not proportional to the original incident. It is proportional to everything that placement could have touched.

The Real Question

The question is not whether your organization will experience candidate fraud. At any meaningful placement volume, friction costs are virtually certain and direct financial losses are probable.

The question is: **where does your organization sit on this continuum, and are your defenses calibrated to your actual risk profile?**

Most organizations absorb Level 1 and 2 costs without measurement. Some experience Level 3 harm without connecting it to fraud. And virtually none have assessed their exposure to Level 4 scenarios – which is precisely where the risk is most asymmetric: low probability, catastrophic impact.

Section 3:

The Hiring Lifecycle

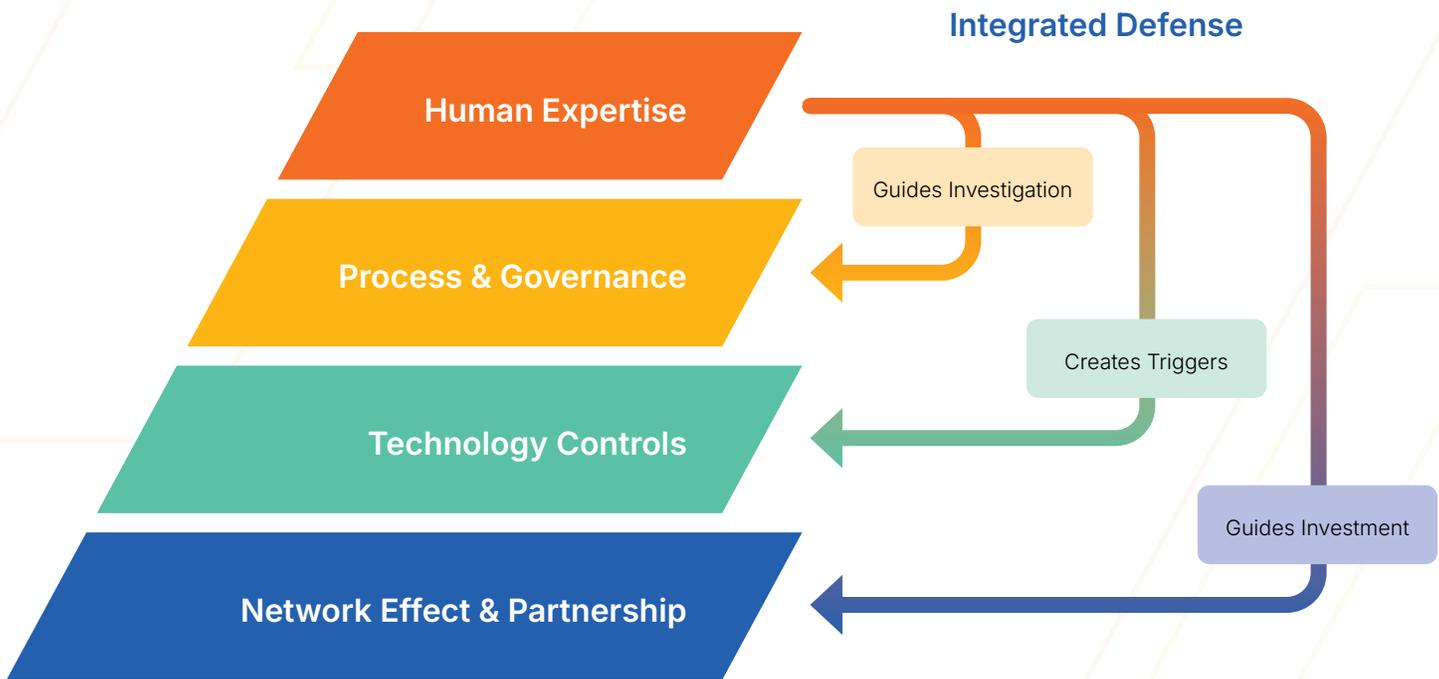
Defense Model

Effective fraud defense requires controls at every stage of the hiring and worker lifecycle, not just at the point of hire. Organizations that concentrate verification in a single phase create predictable gaps that sophisticated actors exploit.

	<i>Pre-Hire</i>	<i>During-Hire</i>	<i>Onboarding</i>	<i>Post-Hire</i>
Technology	<ul style="list-style-type: none"> • Credential verification tools • Resume/profile anomaly detection • Digital identity screening 	<ul style="list-style-type: none"> • Identity-verified video interviews • Monitored technical assessments • Interview recording systems 	<ul style="list-style-type: none"> • ID verification and biometric checks • Device registration controls • Role-based access provisioning 	<ul style="list-style-type: none"> • Login and location monitoring • System activity tracking • Network anomaly detection
Process	<ul style="list-style-type: none"> • Structured credential checks • Standardized resume review • Reference verification 	<ul style="list-style-type: none"> • Multi-step interviews • Live technical demonstrations • Deep-dive experience questions 	<ul style="list-style-type: none"> • Identity confirmation • Security onboarding • Controlled system access setup 	<ul style="list-style-type: none"> • Access reviews • Manager performance check-ins • Monitoring for conflicts or subcontracting
Human Expertise & Cultural Norms	<ul style="list-style-type: none"> • Recruiter fraud awareness • Manager verification mindset • Probe candidate claims 	<ul style="list-style-type: none"> • Technical + behavioral interview panels • Probing follow-up questions • Interviewer fraud training 	<ul style="list-style-type: none"> • Managers validate early work • Encourage disclosure of constraints • Reinforce accountability 	<ul style="list-style-type: none"> • Managers escalate red flags • Recognize fraud as performance risk • Security-aware culture

Within each phase, controls fall into three modalities:

- **Technology:** Automated verification, biometric authentication, behavioral analytics, geolocation, third-party tools and testing
- **Process:** Structured interviews, staged access protocols, cross-functional coordination between HR, Security, and Legal
- **Human Expertise & Cultural Norms:** Interviewer training (identity verification, deepfake recognition, probing questions that defeat coached answers), mandating in-person interviews, fraud indicator awareness, and organizational practices – camera-on policies, in-person interactions, live collaboration – that create continuous verification touchpoints



The third modality is the one most organizations underinvest in – and the one that determines whether the other two actually perform. Technology is easiest to procure, process requires organizational change, but human skills and cultural norms are the fraud-resistant layer that sophisticated operations cannot engineer around. Trained interviewers who can force candidates off rehearsed answers, managers who recognize avoidant behavior patterns, and organizational habits that make sustained deception operationally difficult – these are the capabilities that close the gaps technology and process leave open.

This matters particularly for contingent workforce programs. Organizations managing large contractor populations through multiple staffing vendors and EOR providers face a compounding challenge: fraud detection capability may vary significantly across their supply chain, and the client organization itself may lack visibility into which controls are consistently executed and which degrade under hiring velocity.

Phase 1: Pre-Hire – Building Trusted Pipelines

The strongest fraud defense starts before a requisition opens.

Curate trusted talent networks. Build and maintain relationships with known, vetted talent pools – whether through staffing partners, alumni networks, or communities of practice. Candidates sourced through trusted channels carry lower baseline risk than cold applicants.

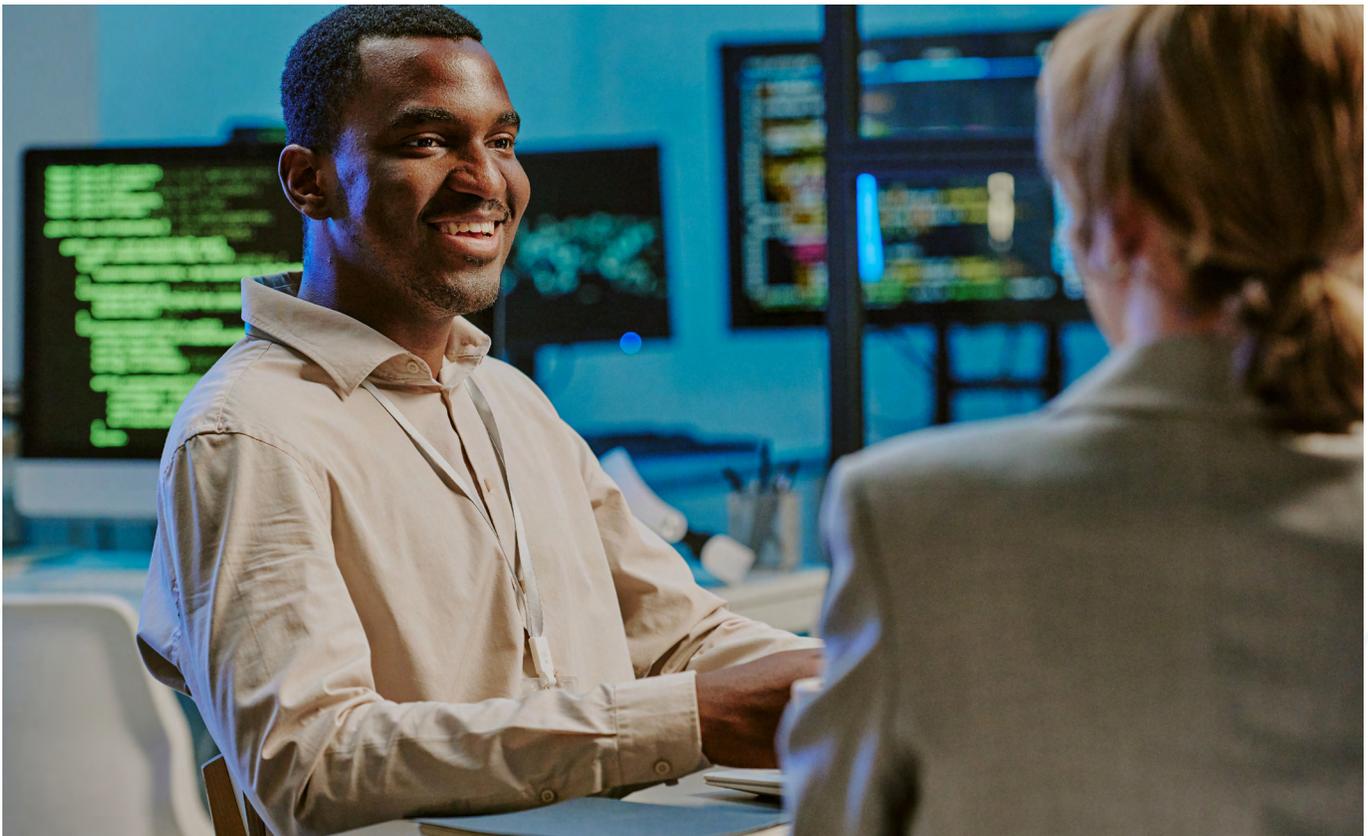
Vet sourcing partners and suppliers. If you use staffing agencies or EOR providers, evaluate their fraud detection capabilities. Your supply chain's weakest verification standard becomes your effective standard. Contractual requirements should specify fraud mitigation obligations, warranties, and indemnification.

Close the direct-source gap. Most organizations have a direct-source funnel where hiring managers identify and engage workers outside vendor channels. Ensure this source is covered by the same fraud controls

as corporate and vendor sources. In EOR/AOR programs, where client managers select workers independently, this gap is often the largest unexamined vulnerability.

Finally, keep standards firm and don't take shortcuts; in many cases, certain parties in the EOR/AOR business may want to skip background checks, ID verification, or even resumes, but these can all act as fraudulent detection points. Prevent gaps from opening at the source. **Define role-specific risk profiles.** Not every role carries the same fraud exposure. Remote positions with system access, roles handling sensitive data, and high-volume contingent programs warrant enhanced controls. Calibrate investment to risk.

Establish AI-use policies before you need them. Define where AI assistance is acceptable in your hiring process and where it is disqualifying – then communicate those boundaries to candidates.



Phase 2: During Hire – Verifying the Candidate in Real Time

The interview and assessment phase is where capability inflation, identity substitution, and proxy interviewing are most likely to succeed – and where real-time controls have the highest detection value.

Mandate at least one live, proctored interaction. Whether in-person or via monitored video, at least one assessment should occur under conditions where the candidate cannot receive external assistance. This single control defeats proxy test-takers and off-screen coaching.

Implement identity verification early. Use selfie-to-government-ID matching before or during the interview process – not after an offer is extended. Deploying identity verification early avoids investing significant time and resources in candidates who may not be who they claim.

Train hiring managers on fraud indicators. Eye movement suggesting off-screen reading, audio artifacts from earpieces, unusual camera

positioning, rehearsed answers that break down under follow-up probing, and reluctance to enable video are observable red flags that untrained interviewers miss. This training is especially critical for client hiring managers in EOR/AOR programs, who are often the only human verification point in the process.

Use technical assessments with biometric confirmation. For technical roles, pair skills assessments with session recording, browser lockdown, or biometric verification that confirms the person completing the test is the person who will show up for work.

Conduct digital profile and IP analysis. Cross-reference the candidate's stated location against their IP address during remote interviews. Audit their digital footprint – LinkedIn history, GitHub contributions, professional community participation – for consistency with claimed experience. Thin or recently created profiles warrant additional scrutiny.

Phase 3: Pre-Start / Onboarding – Closing the Verification Gap

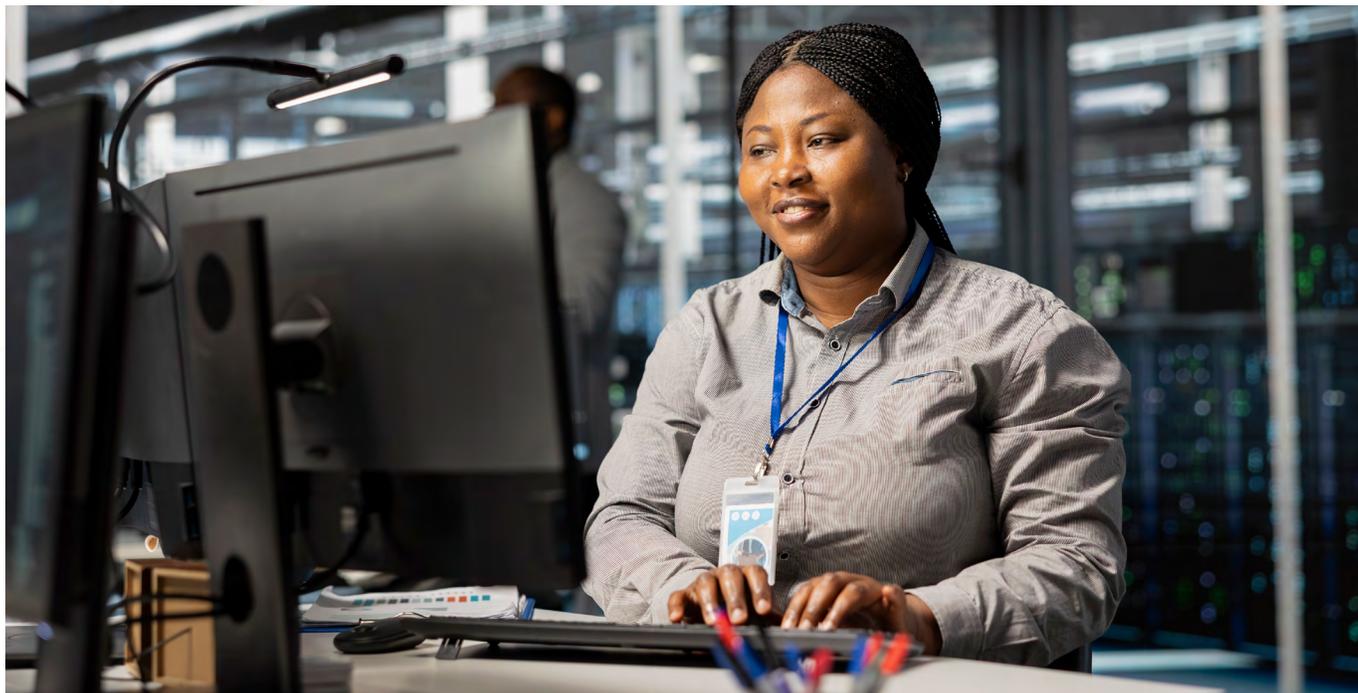
The period between offer acceptance and active work is a critical vulnerability window.

Conduct enhanced reference checks. Contact references identified through independent research – not exclusively those provided by the candidate. Verify the reference's identity and role through company directories or LinkedIn rather than relying on v phone numbers.

Complete background and education verification. Confirm employment history, credential claims, and criminal background through third-party verification services. Cross-reference results against claims made during the hiring process, not just against the application.

Re-verify identity at onboarding. Confirm that the person receiving equipment, credentials, and system access is the same person who was interviewed. Device shipping address should match the candidate's stated location – mismatches are a documented indicator of identity substitution and unauthorized subcontracting. Oftentimes, ID inconsistencies can be detected in supplementing documents such as in IDs, cheques, and tax forms.

Implement staged access protocols. Grant system permissions progressively based on role requirements and demonstrated trustworthiness during the first 90 days. Full access to sensitive systems, code repositories, and client data on day one maximizes damage potential if fraud is detected later.



Phase 4: Post-Hire – Continuous Verification

Several fraud vectors – capability inflation, unauthorized subcontracting, and undisclosed overemployment – only become visible through ongoing monitoring.

Deploy geolocation and IP analysis.

Cross-reference login locations against the worker's stated location and HR records. Fraudulent remote workers frequently generate 'improbable travel' alerts – logins from opposite sides of the country within minutes – when VPN configurations slip.

Monitor network latency patterns. Unusually high or inconsistent latency on video calls and system connections can indicate traffic routing through proxy infrastructure or overseas relay points – a signature of unauthorized subcontracting.

Require periodic in-person interaction. Even for remote roles, scheduled on-site days, team events, or client visits create identity verification touchpoints that proxy workers or unauthorized substitutes cannot navigate.

Use behavioral analytics and activity monitoring. Detect anomalies including unusual login patterns, access to data outside role scope, bulk data transfers, and activity patterns inconsistent with a single individual.

Track performance consistency. A significant and sustained gap between interview or probationary performance and ongoing work quality is a primary indicator of capability inflation. Structured 30/60/90-day performance checkpoints – with tasks that require live, unassisted demonstration – surface these gaps before organizational dependency deepens.

Sustain cultural verification practices. The organizational norms established during onboarding – camera-on expectations, collaborative working sessions, and team interactions – must persist beyond the first weeks. When cultural practices erode over time, so does the ongoing verification pressure they create. Organizations that maintain these practices consistently report earlier detection of behavioral anomalies and avoidant patterns that signal potential fraud.

Integration Across Phases – And Across Partners

No single phase provides complete protection, and no single organization in the supply chain has complete visibility. The model's value comes from reinforcing controls across both dimensions.

Across the lifecycle, each phase backstops the others: pre-hire sourcing quality reduces the volume of fraudulent candidates entering the pipeline; real-time interview controls catch those who get through; onboarding verification closes gaps before access is granted; and post-hire monitoring detects what earlier phases missed.

Across the supply chain, intelligence must flow between partners to be effective. IP and geolocation logs from the EOR provider, interview observations from the client hiring manager, behavioral flags from the staffing vendor, and system access anomalies from the client's IT security team all generate fraud signals – but only if they're shared, correlated, and acted on. Organizations that treat each

partner's verification as an independent silo create the same gap as concentrating controls in a single lifecycle phase: fragments of a pattern that no single party can see in full.

The organizations that execute most effectively across both dimensions share a common characteristic: they operate at sufficient volume to recognize patterns. A hiring manager who sees one coached candidate may not recognize the signal. A staffing vendor placing into a single client may not see the cross-client pattern. An organization processing thousands of placements annually – across multiple clients, industries, and geographies – develops the threat intelligence, behavioral baselines, and pattern-recognition capability that makes every phase and every partner more effective.

This is the human expertise of modality at scale. It cannot be replicated by technology alone, built without operational volume, or achieved without deliberate integration across the supply chain.



Section 4: Assessing Your Security Posture

From awareness to threat risk assessment

In cybersecurity, the discipline that translates awareness into action is the Threat Risk Assessment (TRA): a structured evaluation of what you're exposed to, what controls you have in place, and where the gaps create exploitable vulnerabilities.

Most organizations have some fraud mitigation measures in place but have never pressure-tested them against the current threat environment. Verification procedures designed

five years ago may handle Tier 1 threats while leaving Tier 2 and Tier 3 attack vectors unexamined. Controls may exist on paper but degrade under hiring velocity. The cross-functional response required to detect and contain sophisticated fraud – spanning HR, Security, Legal, and business leadership – may never have been exercised.

The four questions below provide a starting framework.



1. What Is Your Attack Surface?

Your attack surface is defined by your placement volume, role types, sourcing channels, and workforce model. Map it:

- How many contingent placements annually, and in what role types?
- What proportion are remote? Which roles carry system access, sensitive data exposure, or regulatory implications?
- What is your historical fraud detection rate – and how confident are you that it reflects what's actually occurring in your program, rather than only what has been obvious enough to surface?
- Have you measured your Level 1 and 2 costs – the friction losses and direct payroll theft most programs absorb as hiring variability without categorizing them as fraud-related?
- Have you assessed your Level 3 and 4 exposure (operational harm and catastrophic scenarios)?
- How many sourcing channels feed your contingent workforce – and which channels have the weakest verification controls at the point of entry?
- Have you identified which roles in your program have productivity and output expectations calibrated to pre-AI baselines – and whether those expectations create an undetectable capacity gap for AI-proficient workers who are simultaneously employed elsewhere?

That last question represents a new dimension of attack surface that did not exist five years ago. A worker billing full-time hours against a role whose workload was scoped before AI-assisted productivity became the norm may be delivering acceptable output while redirecting significant capacity to a second employer – with no behavioral signal that current monitoring would surface. Mapping your attack surface today requires accounting for this gap explicitly.

Most organizations conducting this assessment for the first time discover they've been absorbing fraud costs they hadn't categorized as such, and their exposure at the higher levels is larger than assumed.



2. What Does Your Incident History Tell You?

Incident review is the foundation of adaptive defense. Every confirmed case, near-miss, and unexplained anomaly generates intelligence that either validates controls or exposes gaps.

Start with what you know: placements where fraud was confirmed, workers terminated for misrepresentation, engagements ended due to unexplained performance gaps. Then examine the ambiguous cases – the contractor who left abruptly at 60 days, the new hire whose performance dropped sharply after probation, the candidate who declined every in-person interaction before disappearing.

Then do something most programs have never done: go back through the last twelve months of early terminations and performance exits and ask whether each one was assessed for fraud indicators before it was closed. In the majority of cases, the answer will be no. The default categorization was bad hire. No fraud indicator checklist was applied. No pattern analysis was conducted across incidents. If that describes your program, the incident history you have is incomplete – and the patterns it might reveal are invisible.

For each incident, three questions matter:

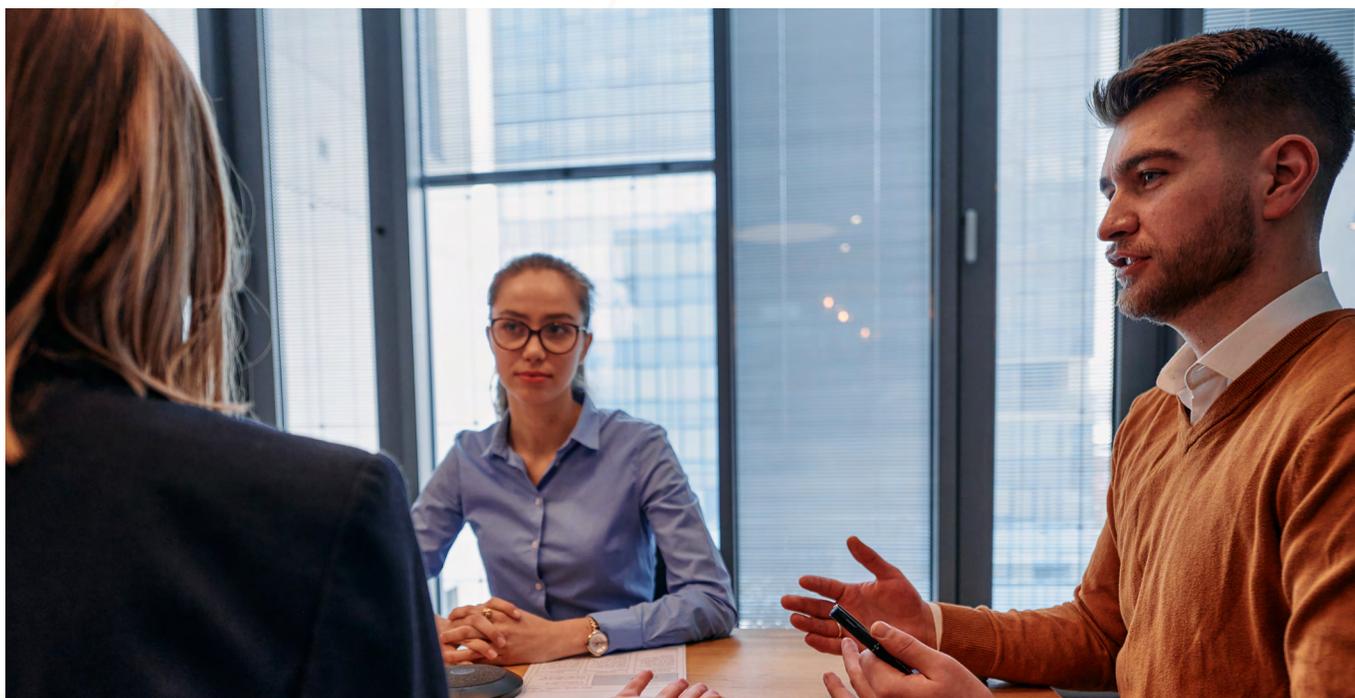
- **Where was it caught – and where could it have been caught earlier?** The gap between actual detection point and earliest possible detection point is a direct measure of control effectiveness.
- **What was the total cost – including costs you didn't categorize as fraud-related?** Re-hiring, project delays, client damage, investigation time, and security remediation are frequently absorbed into general operating budgets rather than attributed to the incident that caused them.
- **Was the intelligence shared or siloed?** If incident knowledge stays with the hiring manager who experienced it, the organization learns nothing. Pattern recognition requires systematically captured data.

Converting exit processes into fraud intelligence

Standard offboarding rarely includes a structured assessment of whether fraud may have been a contributing factor. When a contractor is terminated for performance, the default categorization is 'bad hire' – not 'potential fraud indicator.' Updating exit processes to include a short fraud-indicator checklist – identity discrepancies, unexplained performance gaps between interview and execution, digital profile anomalies, location or access irregularities – converts routine HR processes into an intelligence source. Over time, this data reveals which vectors are active in your environment and where your lifecycle controls are failing.

One important addition to that checklist: assess whether the worker's output and availability patterns were consistent with full-time singular engagement, or whether the pattern – burst activity, inconsistent availability, context-switching errors, work quality that varied sharply by task type – is consistent with divided attention across multiple concurrent roles. This will not produce a confirmed finding in most cases. But over time, across multiple exits, it reveals whether overemployment is an active vector in your program and which role types are most exposed.

This single process change costs nothing to implement and generates the incident data foundation that every other improvement depends on.



3. Where are the Vulnerabilities?

Map your current controls against the four lifecycle phases and identify where coverage decays:

- **Pre-hire:** Are you sourcing from trusted, vetted channels – or accepting the full risk of open inbound pipelines? Have you evaluated your staffing suppliers' and EOR providers' fraud detection capabilities?
- **During hire:** Are identity verification, live proctored assessment, and digital profile analysis consistently executed – or do they get skipped under time-to-fill pressure?
- **Onboarding:** Are background checks, independent reference verification, and identity re-confirmation completing before full system access is granted?
- **Post-hire:** Do you have continuous monitoring – geolocation analysis, latency patterns, behavioral analytics, performance tracking – or does verification end at onboarding?

The typical finding is coverage decay across the lifecycle: reasonable controls during hire, inconsistent execution at onboarding, and minimal post-hire monitoring. This is strong perimeter security with no internal threat detection – exactly the gap sophisticated operations are designed to exploit.

One gap that current monitoring frameworks do not fully address: the worker with a verified identity and a fabricated employment history who performs adequately and generates no behavioral anomalies. Post-hire monitoring as currently practiced catches performance degradation, access anomalies, and behavioral red flags. It does not catch the worker who bills accurately, delivers acceptable output, and simply is not who their resume represents them to be. That exposure sits outside what behavioral monitoring can surface – and requires the employment history verification controls described in Section 3 to have been executed correctly before placement, because there is no reliable post-hire detection mechanism for it.

4. Can You Close the Gaps Internally - or Do You Need to Source Capabilities?

Cybersecurity teams face a well-understood build-versus-buy decision: some capabilities make sense to develop internally, while others – particularly threat intelligence, continuous monitoring, and incident response – are more effective when sourced from specialists operating at a scale that generates better data.

- The same decision applies here. Fraud mitigation requires verification technology, consistently executed processes across every hiring phase, investigative expertise, and continuously updated threat intelligence. The questions that clarify the decision:
- What is the current fraud attempt rate in your industry and role types – and do you have access to cross-industry data to benchmark against?
- How do your verification procedures compare to organizations with similar placement volumes and risk profiles?
- When were your controls last stress-tested against current Tier 2 and Tier 3 tactics – including AI-assisted interview fraud and organized overemployment?
- Do your hiring managers – including client managers selecting contractors in EOR and AOR programs – know what a coached candidate looks like, and do they have clear escalation procedures when something feels wrong?
- Can your team distinguish an organized fraud operation from a nervous legitimate candidate on a bad connection?
- Is your placement volume high enough that verification is executed consistently across every hire – or do controls degrade under time-to-fill pressure?
- Do you have dedicated internal capability for fraud pattern analysis across your program – or does incident knowledge stay siloed with the hiring manager who experienced it?
- Is your post-hire monitoring sufficient to detect the vectors that pre-hire controls miss – or does your verification effectively end at onboarding?

If you can't answer these with confidence, that's a vulnerability worth understanding before it's exploited.



Section 5:

Fraud Posture

Self-Assessment

Use this framework to evaluate your organization's current vulnerability to candidate fraud across the hiring lifecycle. This is a self-administered version of the Threat Risk Assessment described in Section 4 – designed to surface gaps quickly and prioritize where to focus. Answer each question honestly against current practice, not documented policy. Controls that exist on paper but degrade under hiring pressure are not operational controls.

Pre-Hire Controls

- Do you source candidates through trusted, vetted channels – or rely primarily on open inbound pipelines where the quality and integrity of the candidate pool is unknown?
- Have you evaluated your staffing suppliers' and EOR providers' fraud detection capabilities? Do contracts include specific fraud mitigation requirements, warranties, and indemnification?
- Have you defined role-specific risk profiles that calibrate verification intensity to the fraud exposure of each position – accounting for system access, data sensitivity, and remote work status?
- Do your direct-source channels – where hiring managers identify and engage workers outside vendor programs – receive the same fraud controls as vendor-sourced candidates, or does direct sourcing create an uncontrolled entry point?
- Do you have a documented AI-use policy that defines where AI assistance is acceptable in your hiring process and where it is disqualifying – and is that policy communicated to candidates before assessments begin?
- Have you assessed which roles in your program have workload expectations calibrated to pre-AI productivity baselines – creating potential capacity gaps that an AI-proficient worker could exploit while maintaining parallel employment?

During-Hire Controls

- Are interviews conducted with at least one live, proctored interaction where the candidate cannot receive external assistance – human or AI?
- Deploy identity proofing – document authentication, biometric matching, and identity corroboration (selfie-to-government-ID matching) before or during the interview process, rather than after an offer is extended.
- Are technical assessments conducted live by internal experts with biometric confirmation, or through unmonitored take-home exercises?
- Do you conduct digital profile and IP analysis – cross-referencing stated location against interview IP address, and auditing professional presence for consistency with claimed experience?
- Are hiring managers trained to recognize fraud indicators –including the coached candidate, the AI-assisted interviewee, and the resume that reads more fluently than the person presents – including client hiring managers in EOR/AOR programs who may be the only human verification point?
- Do you verify credentials directly with issuing institutions, or rely on candidate-provided documentation?

Onboarding Controls

- Are references contacted through independently verified channels rather than candidate-supplied phone numbers?
- Do background and education checks complete before full system access is granted?
- Does onboarding include identity re-verification – confirming the person receiving equipment and credentials is the person who was interviewed?
- Does the device shipping address match the candidate's stated location?
- Do you implement staged access protocols – limiting system permissions during an initial probationary period?
- Is there a structured 30/60/90-day assessment that evaluates whether demonstrated capability matches interview performance?

Post-Hire Monitoring

- Do you have geolocation and IP analysis capabilities that cross-reference login locations against the worker's stated location?
- Can your systems detect simultaneous access from multiple geographic locations or improbable travel patterns?
- Do you monitor network latency patterns on video calls and system connections for indicators of traffic routing through proxy infrastructure?
- Do you require periodic in-person interaction – even for remote roles – creating identity verification touchpoints?
- Do you monitor for behavioral anomalies: work patterns inconsistent with a single individual, communication style shifts, unusual data access, bulk transfers, or burst-and-gap activity patterns that suggest divided attention across multiple concurrent roles?

- Are managers trained to recognize the specific signature of overemployment – not just performance degradation, but inconsistent availability, context-switching errors, and output that is adequate in quality but limited in volume relative to a fully engaged worker?

Offboarding and Incident Capture

- When workers are terminated for performance – particularly within the first 6 months – does your exit process assess whether fraud may have been a contributing factor?
- Do you screen for fraud indicators at termination: identity discrepancies, unexplained performance gaps, location or access irregularities, digital profile anomalies?
- Does your exit assessment include a review of whether the worker's output volume, availability patterns, and context-switching errors were consistent with full-time singular engagement – or whether the pattern is consistent with divided attention across concurrent roles?
- Are confirmed and suspected fraud incidents documented in a centralized system – or does intelligence stay siloed?
- Do you conduct post-incident analysis to identify at which lifecycle phase the fraud could have been detected earlier?
- Is incident data reviewed periodically to identify patterns across roles, sourcing channels, geographies, or suppliers- not just assessed case by case?
- Are lessons from incident review fed back into hiring process updates, or does each incident get closed without organizational learning?

Organizational Readiness

- Is fraud mitigation treated as an enterprise risk management priority with cross-functional ownership across procurement, HR, Security, and Legal – or does it sit solely with one function that lacks the authority or visibility to act across the others?
- Do you have a documented incident response plan for suspected candidate fraud involving HR, Security, Legal, and business leadership?
- Can you quantify your current fraud exposure across the cost continuum– Level 1 through Level 4 – or are you managing a risk you have never formally measured?
- When was the last time your controls were stress-tested against current Tier 2 and Tier 3 tactics?
- Do you know whether your program is currently experiencing fraud at any level – or is your answer to that question entirely dependent on whether something obvious has surfaced?

Scoring Your Posture

Score	What it means
<i>Mostly "Yes"</i>	You have foundational defenses in place. Focus on consistent execution under hiring pressure and verify that your controls address Tier 2 and Tier 3 threats – not just resume padding.
<i>Mixed results</i>	You have exploitable gaps. Use the Threat Risk Assessment framework in Section 4 to identify which gaps create the highest exposure given your industry, role types, and placement volume – then determine whether closing them requires internal investment, external capability, or both.
<i>Mostly "No"</i>	You have significant unexamined exposure. Start with the attack surface mapping in Section 4: understand your placement volume, role risk profiles, and where you sit on the cost continuum. That assessment is the prerequisite for every other decision.





Closing

The threat landscape for candidate fraud continues to evolve – in sophistication, in scale, and in the tools available to both attackers and defenders. The frameworks in this paper – the six-mechanism taxonomy, the cost continuum, the hiring lifecycle defense model, and the fraud risk posture self-assessment – provide a foundation for assessing and addressing that risk.

Three principles should guide what comes next. First, treat candidate fraud as an enterprise risk management challenge with cross-functional ownership – not an HR inconvenience handled by a single team. Second, assess your controls against the current threat environment, not the one that existed when your verification procedures were designed. Third, recognize that the most effective defenses combine technology, process, and human expertise across the full hiring lifecycle and across every partner in your supply chain – and that the human expertise layer, the hardest to build, is the one that determines whether the other two perform.

The organizations best positioned to manage this risk are those that treat fraud mitigation not as a static set of controls, but as a capability that adapts as fast as the threats it faces – and that recognize the unique concentration of risk in their contingent workforce programs.

To discuss your organization's fraud risk posture, [request a confidential Threat Risk Assessment.](#)