

REBUILDING TRUST IN THE AGE OF AI:

The 2026 Workforce Impersonation Report

How AI-enabled impersonation is redefining identity security and shaping the future of enterprise trust.



Table of Contents

INTRODUCTION	03
<hr/>	
SIX WORKFORCE IMPERSONATION TRENDS	04
DEEPMACHINES & GENERATIVE AI	05
HELPDESK SOCIAL ENGINEERING	06
MFA BYPASS & INTERCEPTION	07
PHISHING & VISHING	08
NORTH KOREAN IT WORKERS	09
AGENTIC AI MISUSE	10
<hr/>	
UNDERSTANDING THE IDENTITY ASSURANCE GAP	11
<hr/>	
SOLUTIONS TO WORKFORCE IMPERSONATION	12
<hr/>	
COMBATING WORKFORCE IMPERSONATION	15



Introduction



Trust itself is now a primary target. Enterprise security strategies must adapt.

In 2025, [organizations reported](#) a growing number of breaches that began with someone pretending to be a legitimate member of the workforce. Nearly every major breach now carries an element of impersonation. Across industries, attackers are impersonating trusted insiders to access systems, people, and data, and legacy trust models are failing.


What once required special skills and resources is now available to anyone with access to AI tools like ChatGPT or Sora. Generative AI has fully blurred the line between real humans and deepfake clones, making it easier than ever for bad actors to convincingly impersonate others. We can no longer trust what we see, hear, or read, or even that someone is a real human.

In [healthcare](#), attackers pose as clinicians or hospital staff to request password or MFA resets, gaining access to internal systems and patient records. In higher education, fraudsters are enrolling [ghost students](#) to access academic resources and steal financial aid. Across [SaaS](#) and [technology](#) companies, bad actors pretend to be legitimate job candidates, current employees, or outside contractors so they can get hired or socially engineer an account recovery. [Transportation and logistics](#) organizations report bad actors disguising themselves as drivers, dispatchers, or vendor contacts to infiltrate distributed networks and redirect operations or payments.

Taken together, these incidents show that traditional device and credential authentication can't keep pace with generative AI and deepfake driven impersonation. To protect their companies, IT and security teams will need to look for systems, technologies, and processes that verify people.



Six Workforce Impersonation Trends



Six Workforce Impersonation Trends Shaping Enterprise Security Strategies in 2026

Adversaries have learned to look, sound, and act legitimate to bypass traditional defenses.

Deepfakes make it easy to steal anyone's [voice or likeness](#). Social engineering scripts are being upgraded with [generative AI](#). MFA bypasses and phishing proxies have evolved into precision tools for scalable session hijacking. Entire nation states are engaged in massive hiring fraud schemes. And autonomous [AI agents](#) are introducing an entirely new class of impersonation risk.

Together, these threats reveal how impersonation has become a linking thread across a wide range of attacks. This reveals a shift from credential-based compromise to identity-layer compromise, where trust itself is the target. Understanding how these threats are evolving is the first step toward restoring confidence in the people, processes, and systems we rely on.

ACCORDING TO GARTNER:

"As adoption accelerates, attacks leveraging GenAI for phishing, deepfakes and social engineering have become mainstream, while other threats – such as attacks on GenAI application infrastructure and prompt-based manipulations – are emerging and gaining traction."

AKIF KHAN, VP ANALYST AT GARTNER

GARTNER PRESS RELEASE, "GARTNER SURVEY REVEALS GENAI ATTACKS ARE ON THE RISE", GARTNER – SEPTEMBER 22, 2025
GARTNER IS A TRADEMARK OF GARTNER, INC. AND/OR ITS AFFILIATES.




Deepfakes & Generative AI



Threat Summary

Deepfakes are synthetic audio, video, and image content created or manipulated by generative AI. What once required technical skills and high-quality source material became mainstream in 2025 with rapid improvements in accessible deepfake tools. Bad actors can now combine text-based AI tools like ChatGPT with AI video and photo generators like Sora 2 to convincingly impersonate others in dynamic, living scenarios that are virtually indistinguishable from reality.

Deepfakes make it substantially easier for bad actors to carry out their impersonation attacks, including helpdesk social engineering and hiring fraud. According to [Gartner](#), 62% of organizations have experienced a deepfake attack in the past 12 months. IT and security teams should think in terms of how a bad actor might use deepfakes (and generative AI more generally) to impersonate someone in a particular communications channel or identity verification process.

EXAMPLE INCIDENT	2026 PREDICTION
<p>In early 2025, a Singaporean company almost lost \$500,000 when a finance director was contacted via WhatsApp by scammers impersonating the CFO. The director joined a video conference where imposters used deepfakes to impersonate senior executives. The incident came a year after British engineering firm Arup lost \$25 million to a scheme wherein a finance employee joined a video conference with what appeared to be the company's CFO and senior colleagues. Investigators later confirmed that every participant on the call was deepfaked.</p>	<p></p> <p></p> <p>Criminals will commercialize "Deepfakes as a Service" (DaaS) kits, further lowering the barrier to entry. These kits will include advanced deepfake injection capabilities designed to trick identity verification systems by inserting false media directly into the data stream; many consumer-grade identity verification systems prove unable to reliably detect injected deepfakes. The ensuing fallout will result in a wholesale shift in enterprise trust models away from static visual checks and AI-based deepfake detection towards continuous, hardware-attested identity verification.</p>



Helpdesk Social Engineering



Threat Summary

Helpdesk social engineering attacks target one of the most vital human interfaces in an organization: the IT support desk. Attackers pose as employees or contractors, using stolen credentials, publicly available personal information, and deepfake voice clones to trick helpdesk staff into resetting a victim's password or MFA. After gaining access to the victim's account, the bad actor performs reconnaissance, escalates privileges, steals data, and deploys ransomware.

In 2025, led by [Scattered Spider](#), breaches traced to helpdesk social engineering saw a resurgence which will continue in 2026. Employee helpdesks and customer support centers are particularly vulnerable to social engineering because agents don't have an effective way to verify the person on the other end of a phone call, chat, or support ticket. The past few years have seen a proliferation of helpdesk verification tools, but the assurance they can provide varies widely.

EXAMPLE INCIDENT	2026 PREDICTION
<p>In early 2025, UK retailers Marks & Spencer and the Co-op Group experienced major disruptions after attackers socially engineered helpdesk representatives at a third-party service provider into resetting credentials for privileged accounts. The attacks, attributed to Scattered Spider, were estimated to cost nearly \$600 million and followed similar attacks against Harrods, Transport for London (TfL) in 2024, and MGM Resorts International and Caesars Entertainment in 2023.</p>	<p>In 2026, deepfake impersonation (voice and video) will become a standard tactic in helpdesk social engineering playbooks. Solutions which rely on traditional authentication factors to verify users will fail to stop these attacks, leading organizations to equip their helpdesks with identity verification tools instead. Attackers will target more managed service providers (MSPs) whose helpdesks grant access to multiple downstream clients, pressuring MSPs to increase security.</p>



MFA Bypass & Interception



Threat Summary

Multi-factor authentication (MFA) remains a critical, baseline security control, but adversaries have developed numerous reliable workarounds. A survey conducted by [Portnox](#) found that 96% of CISOs believe MFA “can’t keep up with today’s threat landscape.” Common tactics observed in 2025 include push notification (“push fatigue”) attacks; SIM swap attacks that move a victim’s phone number to an attacker’s device; and adversary-in-the-middle (AitM) techniques which steal authenticated session tokens.

Other attacks bypass MFA entirely by targeting its recovery process. Many implementations of MFA, even where they use phishing-resistant MFA, downgrade to a password or one-time passcode when someone clicks “Can’t access my MFA.” As passwordless MFA adoption continues in 2026, more organizations will implement additional security around MFA recovery.

EXAMPLE INCIDENT	2026 PREDICTION
<p>In early 2025, researchers uncovered a large-scale AitM campaign to intercept MFA challenges and steal valid Microsoft 365 session tokens. Later in the year, investigators found that threat actors were quietly harvesting OAuth tokens through malware and browser-based session hijacking, letting them replay authenticated sessions. These incidents show how attackers have industrialized a range of MFA exploitations to impersonate legitimate, authenticated users.</p>	<p>In 2026, successful MFA downgrade attacks which target humans and processes to bypass phishing-resistant MFA and passwordless factors will lead more organizations to implement a policy of using identity verification to upgrade security during MFA resets. Adversaries will combine MFA interception with MFA bypass in multi-stage campaigns to defeat mixed deployments. Hardening of the processes surrounding MFA will be key. Insurers and even regulators will increasingly require phishing-resistant MFA and upgraded recovery security.</p>



Phishing & Vishing



Threat Summary

Phishing and vishing (voice phishing) is the act of sending messages claiming to be from a trusted individual or company. The purpose of a phishing attack is typically to trick a victim into divulging sensitive information like login credentials, and/or to facilitate MFA interception. In 2025, phishing attacks remained one of the most common forms of impersonation. And now, attackers are using generative AI and deepfake tools to make their messages more convincing.

2025 saw an increasing shift towards AI-assisted [polymorphic phishing campaigns](#) and a diversification of channels; 34% of [phishing attacks](#) now come through non-email channels. Bad actors are weaponizing generative AI to craft convincing lures which tailor their tone, context, and timing to individual recipients. While defenders continue to focus on awareness training and email filters, attackers are reaching victims through collaboration platforms and social media.

EXAMPLE INCIDENT	2026 PREDICTION
<p>In August 2025, researchers at Group-IB revealed an advanced vishing campaign wherein attackers posed as recruiters from major tech companies, using deepfake voice clones in follow-up calls after phishing emails. Victims were persuaded to download “application forms” that delivered credential-stealing malware. The combination of believable correspondence and synthetic voice made the scheme nearly indistinguishable from a legitimate hiring process.</p>	<p>In 2026, scalable, AI-personalized spear phishing kits will render traditional phishing cues nearly useless. Criminals will operationalize hybrid phishing campaigns using AI-generated text, voice, and video messages to launch individually-personalized attacks en masse. Organizations begin implementing policies which require out-of-band authentication, using peer-to-peer verification solutions which equip employees with a way to easily authenticate the person they’re talking to.</p>



North Korean IT Workers



Threat Summary

Thousands of operatives working for the Democratic People's Republic of Korea (DPRK) have infiltrated companies around the world through remote IT jobs. Their goal is to collect paychecks, steal source code, and exfiltrate data. They frequently use generative AI and deepfakes to mask their true identities, combining stolen identities with fake IDs and even deepfake video filters.

DPRK IT worker schemes are a form of hiring fraud where bad actors secure jobs at target organizations by posing as legitimate job candidates. Traditional hiring checks like background checks, video interviews, and I-9 validation are proving ineffective. The disconnect between HR teams responsible for hiring distributed workforces, and IT teams responsible for onboarding them, creates a set of [inherited risks](#) between departments.

EXAMPLE INCIDENT	2026 PREDICTION
<p>In July 2025, an American woman was sentenced to over 8 years in prison for her role operating a laptop farm on behalf of North Korean IT workers. The particular ring had reportedly breached more than 300 American companies and government agencies, but threat intelligence researchers have confirmed thousands more infiltrations around the world, including the Fortune 500. It's estimated that fake IT workers have funneled up to \$1 billion to the DPRK regime.</p>	<p>In 2026, regulatory crackdowns and mandatory identity verification requirements will reshape hiring security, but adversaries will adapt. The U.S., U.K., and E.U. will begin penalizing firms found employing workers from the DPRK and other sanctioned entities. Other state-aligned and financially motivated actors will emulate DPRK tactics, building "fake IT worker" networks for profit or espionage. Enterprise HR departments and freelancer platforms will adopt policies of mandatory identity verification for job applicants, with re-verification at IT onboarding.</p>



Agentic AI Misuse



Threat Summary

Agentic AI are autonomous software agents that can reason, plan, and act without direct human prompts. In 2025, AI agents became a first-class identity, gaining more autonomy and wider access to enterprise applications and other AI agents. But the transition to fully agentic workflows reveals major gaps in how enterprises define, manage, and protect AI identities.


For security teams, agentic AI is a serious new threat. Once an agent has access and autonomy, it can be hijacked or misused just like any other user account. If an attacker compromises an agent identity or manipulates its decision logic. They can initiate legitimate-looking actions, from data exports to software deployments, that bypass human oversight. Traditional IAM systems, designed for static human or service accounts, rarely capture which human authorized which agentic action, creating new blind spots which attackers are already exploiting.

EXAMPLE INCIDENT



In August 2025, [Anthropic revealed](#) that attackers used its Claude Code tool to automate an entire data-extortion campaign with only minimal human input. Months later, the company [uncovered](#) what appeared to be the first AI-run espionage operation, in which an autonomous agent performed the vast majority of intrusion activity on its own. These incidents demonstrate how bad actors can impersonate trusted users in order to abuse the autonomy and access granted to AI agents, and the importance of verifying the real human behind an AI's actions.

2026 PREDICTION



In 2026, more breaches stemming from over-permissioned AI agents will lead agentic identity governance to become a board-level and compliance priority, akin to supply chain and insider risks. IAM providers will continue to expand their platforms to cover both human and non-human identities under unified trust models. Standard policy will require verified human signatures for risky actions, defined via IAM policy, performed by AI agents on behalf of humans.

WHY WORKFORCE IMPERSONATION WORKS:

Understanding The Identity Assurance Gap

Workforce impersonation works for a simple reason: most authentication methods verify access or ownership, not identity. Standard authentication factors and enterprise trust models assume that whoever can unlock a device, receive a code, or tap a push notification is legitimate. But attackers have countless ways to steal, spoof, or bypass those signals.

The identity assurance gap is the fact that most organizations don't know who is actually behind any given account, device, or action. This gap exists because most authentication factors verify access to a particular device or set of credentials, not an actual person.

Today, companies of every size still struggle to know who they're hiring, onboarding, and giving access. Responsibility for identity assurance is spread across several teams: recruiters screen resumes, HR runs background checks, IT issues devices, security monitors authentication logs, and support teams restore access when something breaks. No one owns the full picture.

These disconnects create predictable openings. During hiring, recruiting teams rely on ill-defined visual and audio cues to spot fraudulent candidates. HR validates a person's ID document, but not that the person is actually holding it. IT ships a laptop to an address HR tells them, and trusts that the person receiving it is legitimate. Helpdesks ask for basic personal info to restore access.

In the end, you have little assurance that the new employee who now has a corporate laptop, security key, and directory credentials is not impersonating someone else in some way.

"Insider threats are evolving faster than most people realize; economic uncertainties and AI tools are combining to create the perfect storm. The reality is that many companies don't actually know who they're hiring and onboarding. In 2026, security teams will focus more on collaborating with HR and IT to implement stronger identity checks and collate signals across departments, resulting in fewer infiltrations. On the other hand, organizations that don't make a concerted effort to strengthen their defenses against hiring fraud will feel the consequences."

CHRIS O'ROURKE, SENIOR MANAGER, CLOUDFORCE ONE - R.E.A.C.T., CLOUDFLARE

Once onboarded, companies use a variety of factors to verify that an employee signing into an app, opening a helpdesk ticket, or resetting another authentication factor is legitimate. But most authentication factors validate that someone has access to a device or credential, not the actual person behind it. Even the stronger authentication factors can still be bypassed through downgrade attacks on the recovery or reset process.

Until organizations verify the person they're hiring and the person behind each session and action, bad actors will continue to exploit the assurance gap inherent in modern identity security.

Solutions to Workforce Impersonation

Impersonation attacks work because most authentication factors check what someone has, not who they are. To stop imposters, organizations must begin verifying the person behind any given account, session, and request. Effective protection requires a layered approach. In 2026, that will mean combining phishing-resistant MFA with workforce identity verification.

Phishing-Resistant MFA

Passwordless authentication factors like device-bound passkeys and physical security keys are widely regarded to be the most secure of the current authentication factors. In 2026, we expect phishing-resistant authentication will become a baseline requirement; all employees will be required to sign in and re-authenticate using phishing-resistant factors.

However, when someone loses or changes their device, they'll need to go through a recovery flow which typically downgrades to a lower assurance factor. If a bad actor can click "Can't access my passkey" and then enter a password or OTP, the entire phishing-resistant system collapses. Ultimately, any authentication factor is only as secure as its recovery process.

The same risk applies to onboarding new users. Allowing employees and contractors to enroll in MFA without proper verification could result in giving legitimate credentials to imposters. As IT departments roll out passwordless authentication, they must account for this identity gap. Building on momentum grown over the past few years, 2026 should see more organizations recognizing these gaps and adopting workforce-grade identity verification solutions, like Nametag, to protect the enrollment and recovery flows that underpin phishing-resistant MFA.

"Scale, efficiency improvements and security are driving businesses to focus on achieving value for their security projects. To deliver the maximum value from their MFA solutions, companies must replace existing, phishable processes for registration and recovery with easy-to-use, self-service and highly secure alternatives."

These foundations will enable enterprises to accelerate their deployments of phishing-resistant authentication rapidly in 2026. The goal is to eliminate phishable authentication methods and identity processes from the enterprise by the end of the year."

DEREK HANSON, FIELD CTO, YUBICO

The Limits of Video Verification

Video calls are commonly recommended as a mitigation for a range of impersonation threats. The approach gained ground in August 2023, when [Okta's chief security officer suggested](#) adding "a visual verification step at the helpdesk." Through 2024 and 2025, security firms and agencies recommended video interviews to combat DPRK IT workers and other hiring fraud. On the surface, video calls are a good idea, and can indeed help organizations spot some imposters.


In reality, video calls create a false sense of security and should not be trusted.

Live video deepfake filters are [available on GitHub](#) and their quality is improving rapidly. The story of a finance employee being fooled into wiring away [\\$25 million](#) is just the tip of the iceberg; one [recent report](#) warned that cases of AI-powered video calls rose 118% in 2024.

Ultimately, organizations should not trust video calls for verifying candidates and employees. If someone is unable or unwilling to complete identity verification, require them to come in-person.

Workforce Identity Verification

To cover the identity assurance gap, organizations are increasingly turning to workforce identity verification (IDV). Identity verification systems ask a person to scan their government-issued photo ID and then take a selfie in order to prove that they really are who they claim to be. When someone has a valid ID with a photo that matches their selfie, they're verified. If the selfie and ID photo don't match, or if the system spots signs of fraud, the verification is rejected or escalated for human review.



Workforce IDV specifically enhances employee and contractor account security by surrounding enrollment, recovery, and step-up authentication processes with a higher level of assurance. To be workforce-grade, identity verification systems must block injection attacks and deepfakes, work with directories and other enterprise identity technology infrastructure, and deploy quickly.


The past year saw a growing number of enterprises deploying workforce IDV to protect helpdesk agents and enable self-service account recovery, creating cost savings and efficiency gains for IT teams. 2026 will see ongoing adoption and expansion of workforce IDV solutions to cover additional use cases and scenarios. Meanwhile, organized attackers will begin collaborating to trick these IDV systems by combining democratized injection attack techniques with ultra-realistic deepfake identity documents, selfies and videos. In this new threat environment, a clear distinction will emerge between workforce-grade and consumer-grade IDV systems.


Purpose-built workforce IDV systems will prove resilient to increasingly sophisticated and commoditized deepfake impersonation kits. Repurposed consumer-grade IDV systems will fall victim, and then be flooded by bad actors capitalizing on vulnerabilities uncovered by their peers.

Where to Deploy Workforce Identity Verification

Workforce identity verification solutions have the most impact when they're applied at specific, high-risk moments across an employee's journey through your organization. These are the points where the risk of impersonation increases, and where high assurance is an absolute necessity.

Workforce Identity Verification Covers The Employee Lifecycle






Hiring and onboarding: Verify that the person you're interviewing, hiring, and onboarding isn't impersonating someone else, before you give them directory credentials or ship them a device.



Device enrollment and replacement: Create a trustworthy link between company-issued devices and the people behind them, limiting potential exposure from lost, stolen, or "borrowed" devices.



Account recoveries and helpdesk tickets: Treat "I'm locked out" as a security event. Verify the actual person requesting a password or MFA reset, not just a piece of information or a device.



Role and privilege changes: Step up to identity verification when someone requests privileged access or moves into a higher-privilege role, ensuring you know who actually holds that access.



High-value approvals: Require identity verification for actions that move money, expose data, or change security posture. Create an auditable record of which verified person approved the action.



Oversight of AI agents: When AI agents can perform actions on behalf of people, workforce identity verification helps ensure that there is a real person accountable for what the agent does.

CONCLUSION

Combating Workforce Impersonation

In 2026, attackers won't rely on sophisticated exploits; they'll rely on looking and sounding like the people your systems already trust. Deepfake-enabled fraud, social engineering, session theft, and hiring fraud points to the same reality: impersonation has become the primary attack vector.

Traditional approaches to identity security that authenticate devices and credentials, but not actual people, leave a widening identity assurance gap that adversaries are actively exploiting.

Closing that gap will require a fundamental shift in how organizations think about workforce identity. Phishing-resistant MFA must become table stakes, but it cannot stand alone. High-risk events — hiring and onboarding, device enrollment and replacement, account recoveries, role and privilege changes, high-value approvals, and oversight of AI agents — all demand higher assurance. That means verifying that the right human is behind the keyboard, phone, or AI agent, not blindly trusting that whoever can click a link, tap a push, or join a call is who they claim to be.

Organizations that adapt will make the shift from periodic identity check, to continuous identity assurance, embedding workforce identity verification into the places where trust matters most.

Looking to close your
own identity assurance gap?

**Visit getnametag.com
to learn more.**

