GetReal



# Deepfake Readiness Benchmark Report:

## Identity Manipulation, Synthetic Content, and the State of Enterprise Preparedness

# Executive Summary



Long anticipated but often viewed as distant, AI-powered impersonation and deepfake threats have now materialized. This report shows they are no longer theoretical. They are a present, active enterprise risk. Organizations face adversaries using fake images, audio, and video to infiltrate hiring processes, impersonate executives, and socially engineer IT help desks in account takeover schemes.

To benchmark enterprise readiness against this growing threat, we surveyed 668 IT, cybersecurity, fraud and risk leaders at enterprises with 1,000 or more employees.

The results show that most organizations have already encountered AI-powered impersonation attacks and recognize the risk, but also reveal an imbalance in detection and defense capabilities.

**Deepfakes and impersonation attacks are no longer hypothetical, but routine.** Eight out of ten responding organizations encounter AI deepfakes or impersonation attempts at least occasionally, including 45% that encounter them frequently.

**Nearly half (41%) of organizations with 1,000 or more employees report having hired and onboarded a fake job candidate or imposter,** demonstrating an acceleration of adversaries infiltrating hiring processes.

**Enterprises acknowledge the risk, but underestimate their vulnerability.** While 85% of IT, cybersecurity, fraud, and risk leaders believe a deepfake incident would result in significant reputational damage, many overestimate their organization's ability to detect deception or believe MFA alone will protect them.

**Enterprises rely on voice and video calls for verification, but apply security inconsistently.** A majority (80%) of respondents use voice and video calls to verify an individual's identity as part of a business process, yet only half (52%) require authentication to stop impersonation attempts on video conference calls.

**Two thirds of enterprises believe their response plans for a deepfake incident are sufficient.** However, those plans rely heavily on employee training, which quickly becomes outdated as GenAI tools produce increasingly realistic images, audio, and video.

**Just over half of enterprises are adapting their identity and access management strategies for GenAI-powered threats.** The top priorities cited, stronger biometric authentication (63%) and multi-factor authentication expansion (62%), are vulnerable to deepfakes and require continuous monitoring to ensure the person behind a credential is who they claim to be.

# Introduction



Cybersecurity incidents and fraud activity involving digital media generated or manipulated with AI tools, from static images to live-stream audio and video, are increasing in both frequency and sophistication. This synthetic media, commonly referred to as “deepfakes,” uses generative AI (GenAI) to depict events that never occurred or mimic a real person’s face or voice in a way that is difficult to distinguish from reality.

While enterprise leaders may associate GenAI-powered deepfakes with high profile examples such as Pope Francis in a puffer jacket or fake backyard surveillance footage of a goat riding a dirt bike, GenAI is also enabling far more consequential impersonation of employees, executives, customers, partners, job applicants, and more. AI-generated content and related identity manipulation create new attack vectors that adversaries already exploit to infiltrate enterprises through remote hiring, remote collaboration, and digitized business processes.

This shift from novelty to operational threat is also reflected in recent industry data. Gartner reports[1] that 62 percent of organizations have experienced a deepfake attack within the last 12 months:

During a video call with an employee

36%

During an audio call with an employee

44%

Against automated face biometrics or identity verification systems

30%

Against automated voice biometrics systems

32%

Deepfakes also contribute to data breaches, as evidenced by IBM findings[2] that AI deepfake attacks contributed to 6 percent of breaches occurring between March 2024 and February 2025.

A threat once considered theoretical has become a real and immediate problem that is already inflicting damage on enterprises today.

To understand how prepared enterprises are for this new threat, we surveyed 668 IT, cybersecurity, fraud, and risk leaders from organizations with more than 1,000 employees across multiple industries. The results provide a snapshot of enterprise readiness and a benchmark for leaders to measure their own preparedness and practices against those of their peers.

The findings also reveal a disconnect. Although leaders generally recognize the significance of the threat, many seem to overestimate their ability to detect or defend it. At the same time, they underestimate how extensively they will need to adapt cybersecurity, identity, and other practices to protect themselves against accelerating AI-powered impersonation attacks.

[1] https://www.gartner.com/en/newsroom/press-releases/2025-09-22-gartner-survey-reveals-generative-artificial-intelligence-attacks-are-on-the-rise
[2] https://www.ibm.com/reports/data-breach

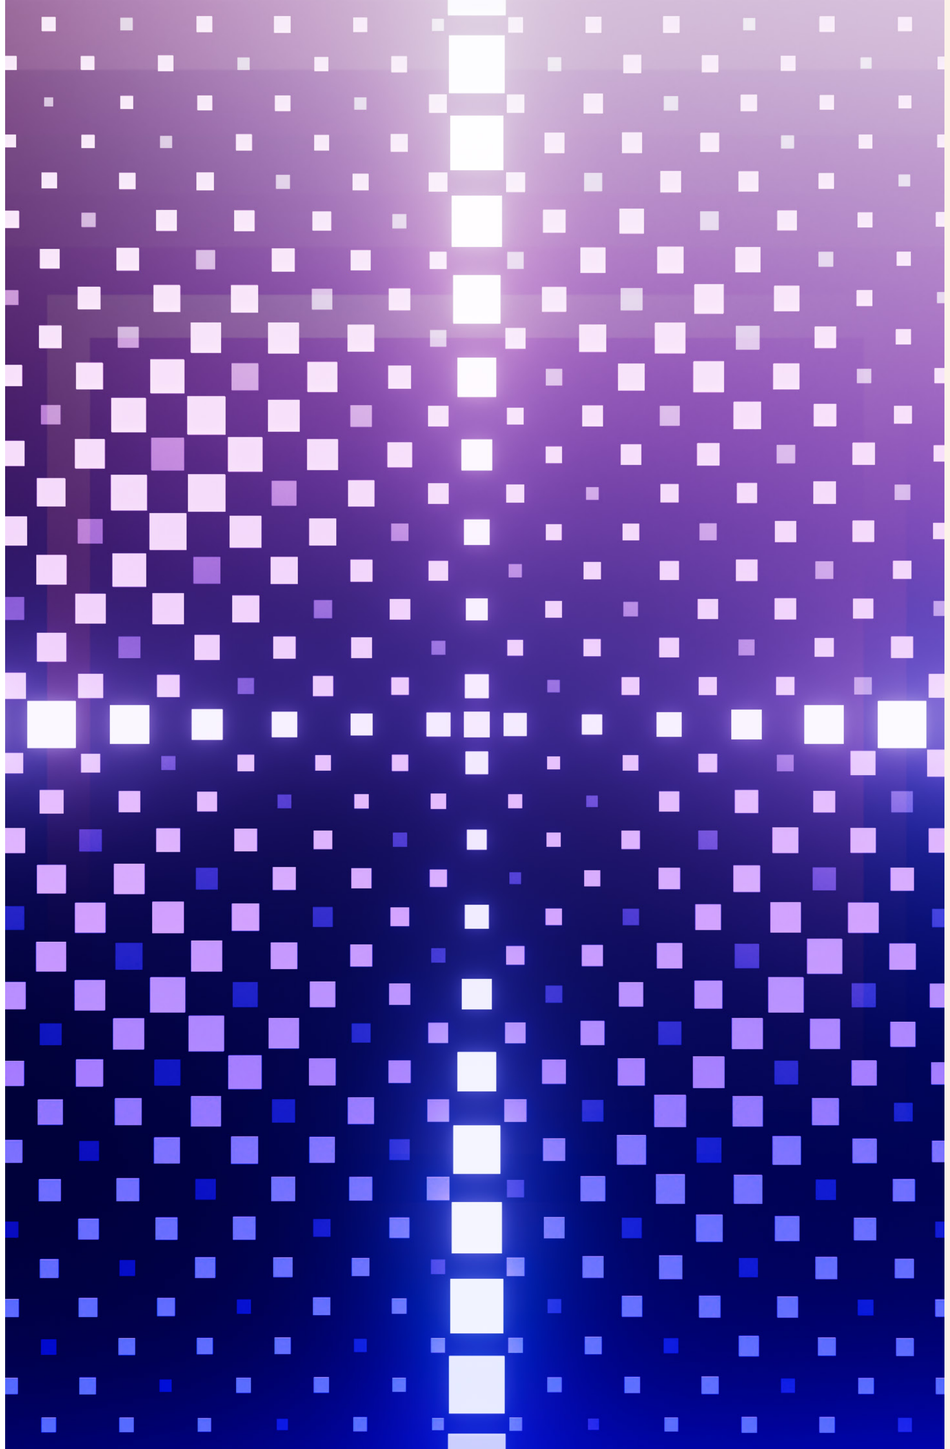

# 1.0/ Deepfake Risk Perception and Exposure Among Enterprises

# 1.1 Enterprise Exposure to Deepfakes

**How frequently does your organization experience deepfakes or AI impersonation attempts?**

Nearly every responding organization (~99%) reported experiencing an AI deepfake or impersonation attempt, though the frequency of incidents varied.

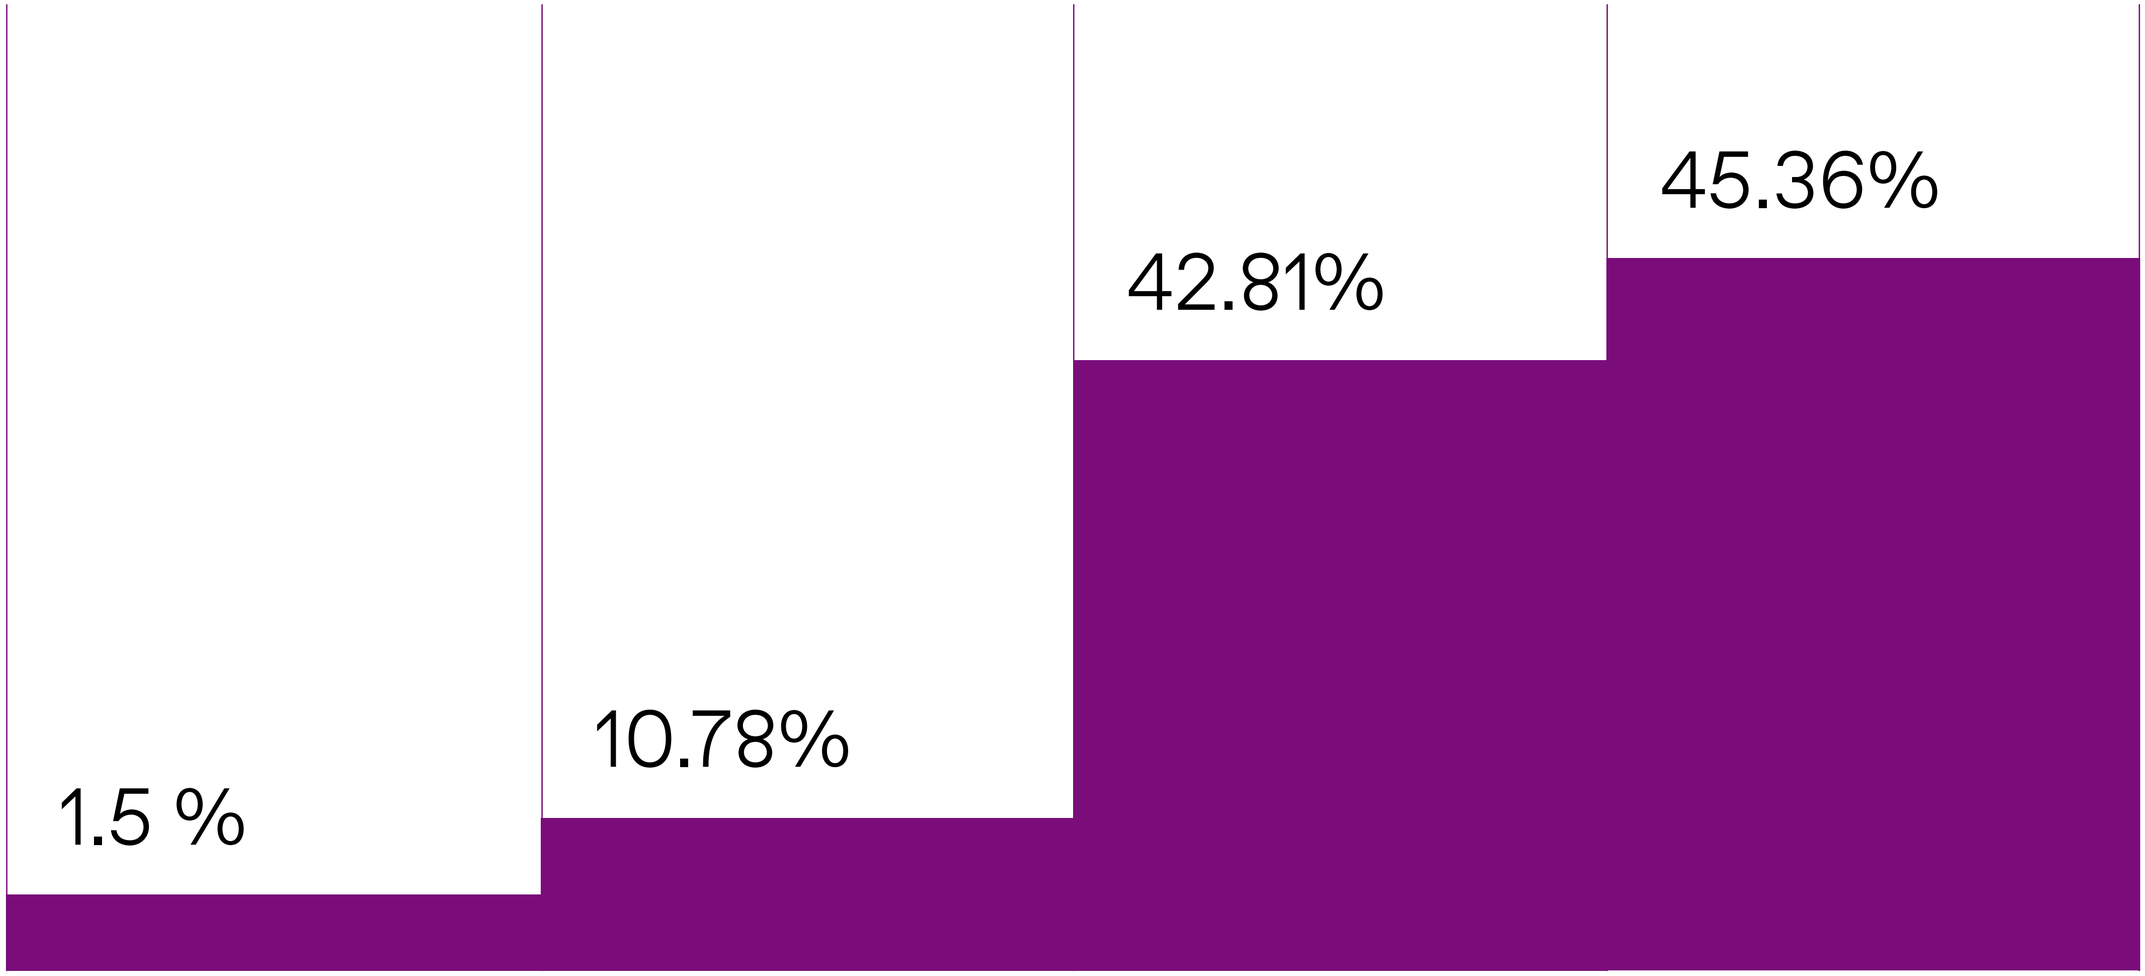

# 1.2 Perception of Deepfake Risk



**How would you rate the level of damage to your organization's reputation as a result of a potential deepfake?**

Approximately 85% of organizations believe a potential deepfake would be significantly or highly damaging to their organization's reputation. Respondents in the financial services industry were slightly more likely to rank a deepfake incident as highly damaging (53% compared to 46% across non-financial services organizations).

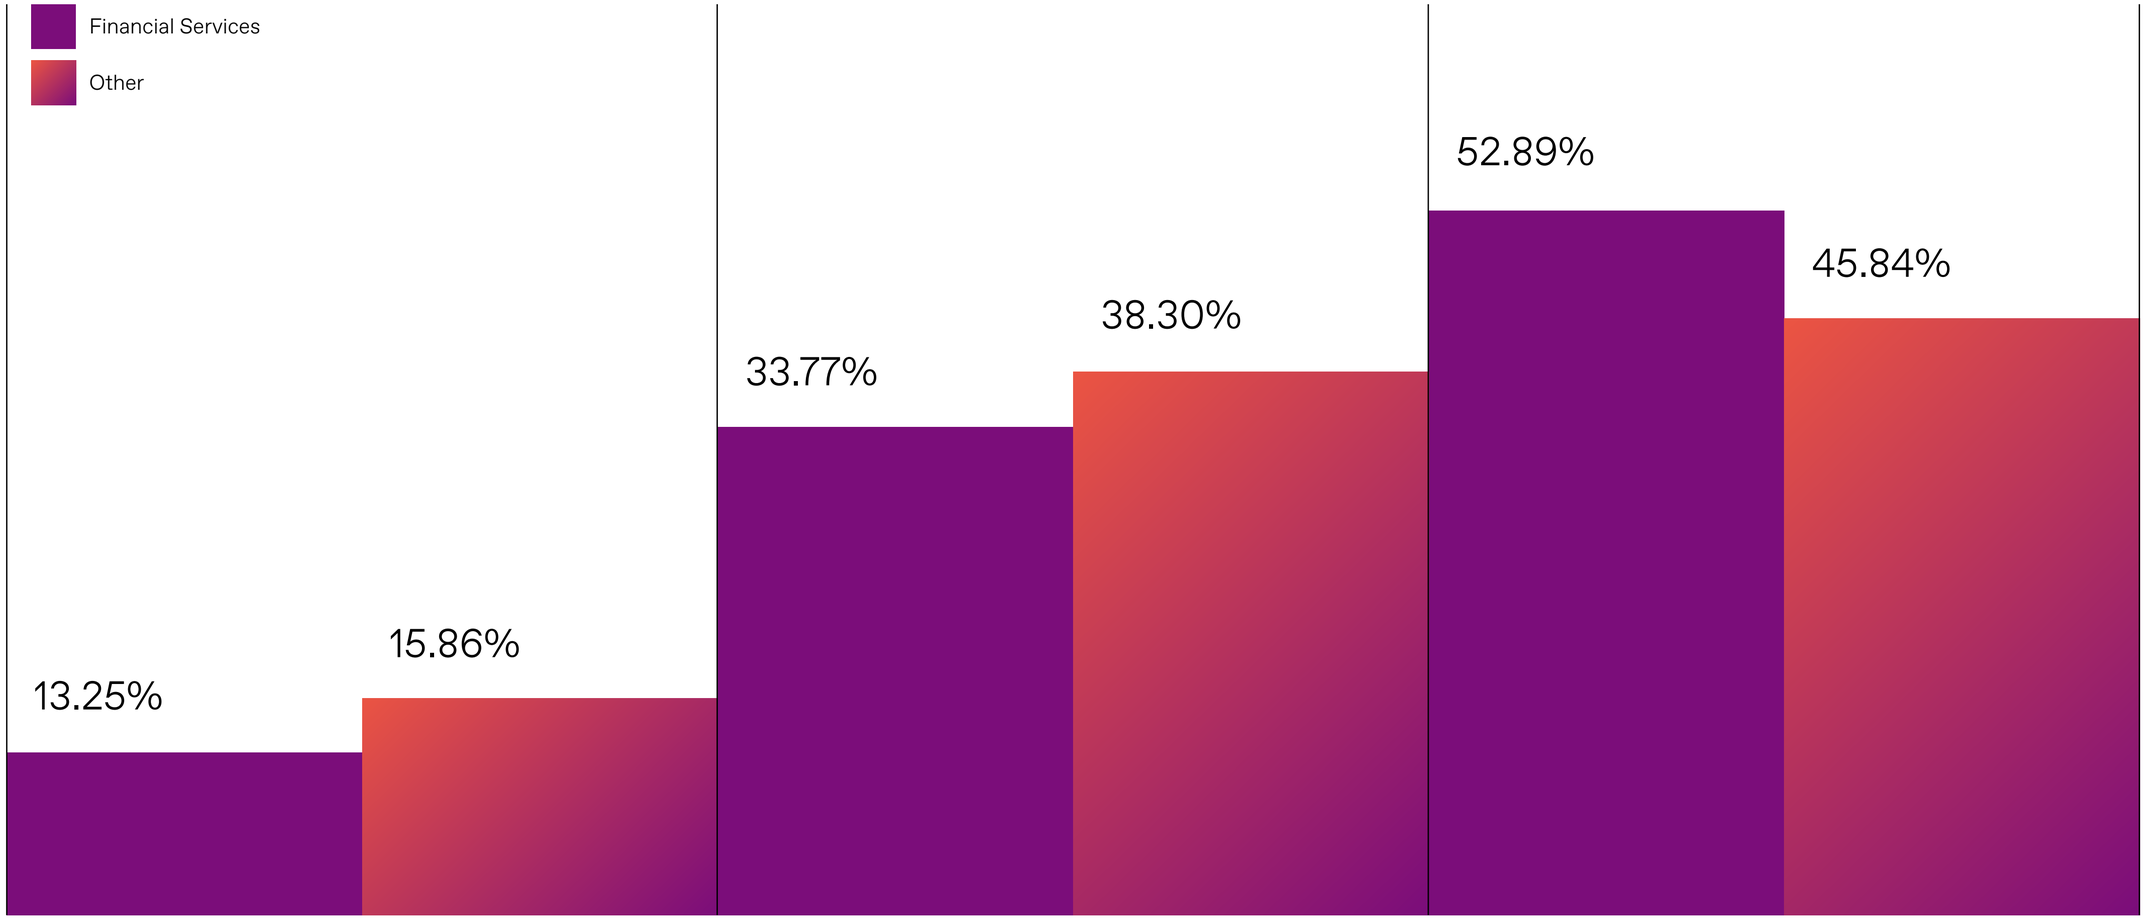

# 1.3 Deepfake Attacks of Greatest Concern to Enterprises

**Which types of deepfake cyberattacks are you currently most concerned about? (Select all that apply)**

No single type of deepfake or AI-powered impersonation scam captured an overwhelming majority of concern from respondents. The top concerns – phishing with deepfake audio or video (56 percent of respondents), voice spoofing in phone scams (56 percent), and impersonation in video conferencing (54 percent) – reflect that enterprises do have some understanding of the risks posed.

The absence of a clear front-runner, however, suggests enterprises view deepfakes as a broader risk rather than focusing on specific attack vectors that have already led to data breaches and financial losses – prime areas where enterprises should begin maturing their deepfake defenses. The low ranking of fake job candidates is particularly troubling, given the well-known rise in nation-state actors, such as DPRK IT workers, successfully infiltrating corporate environments.

While the variety of concerns suggest a recognition of the wide attack surface created by synthetic content and deepfakes, it also suggests a lack of prioritization.

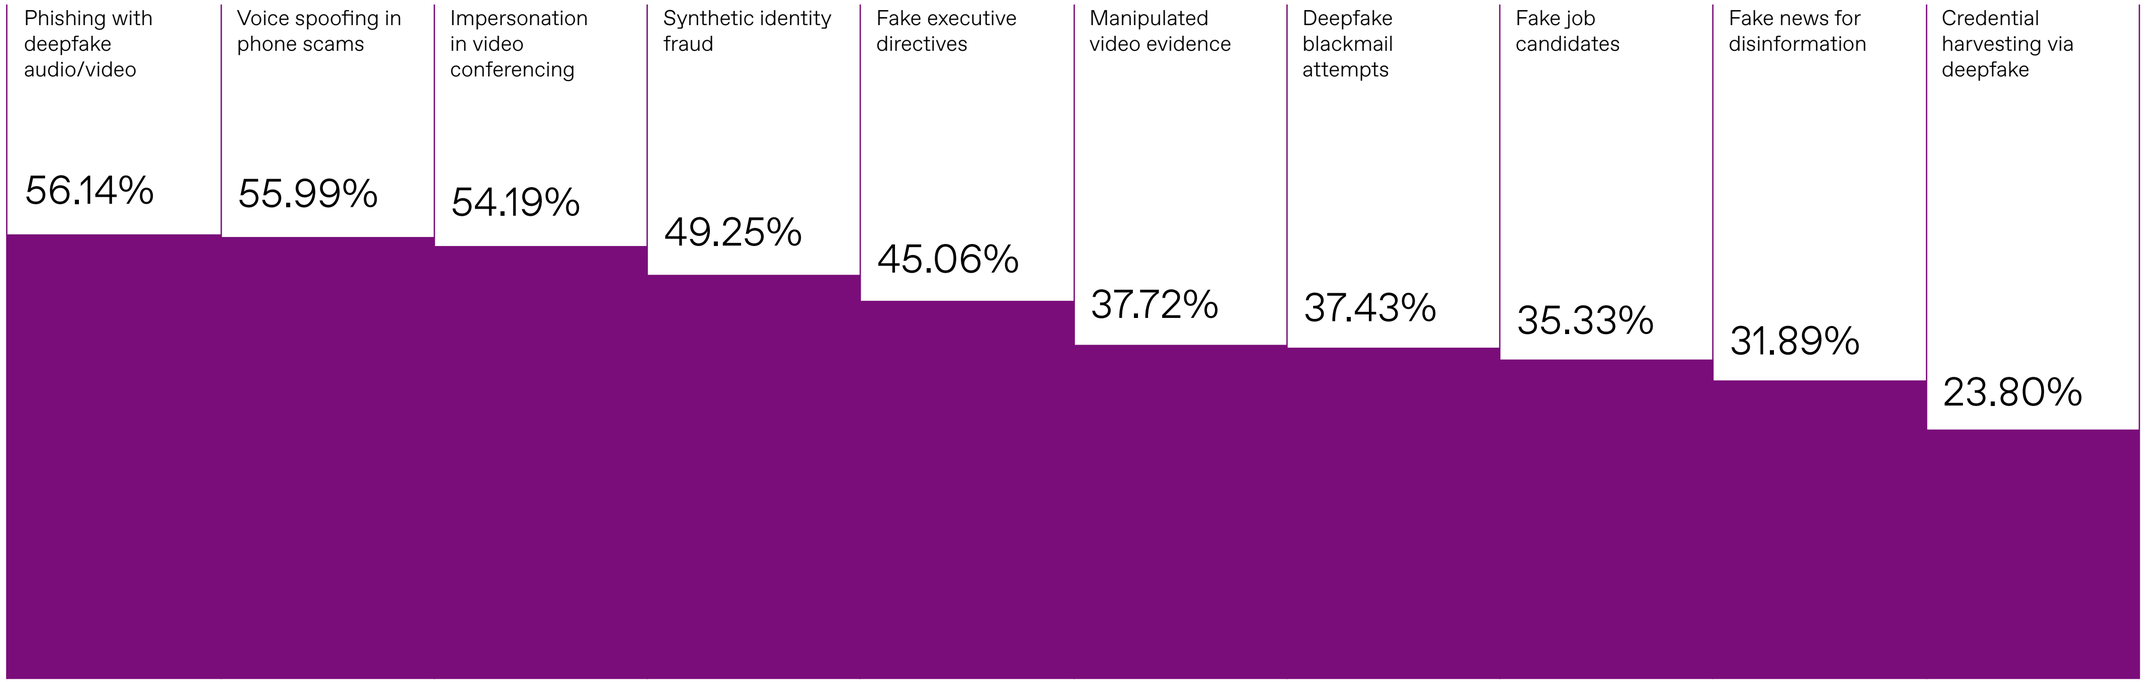

# 1.4 Almost Half of Enterprises Admit to an Imposter on the Payroll

**Was any fake candidate hired and only discovered after starting employment?**

While not surprising in light of reports that nearly every Fortune 500 company has hired a fraudulent IT worker out of North Korea,[3] 41.02 percent of respondents admitted to having hired a fraudulent job candidate that was only discovered after onboarding. This underscores how siloed HR, identity verification, and cybersecurity processes and teams creates openings allowing insider threats to embed within corporate systems.

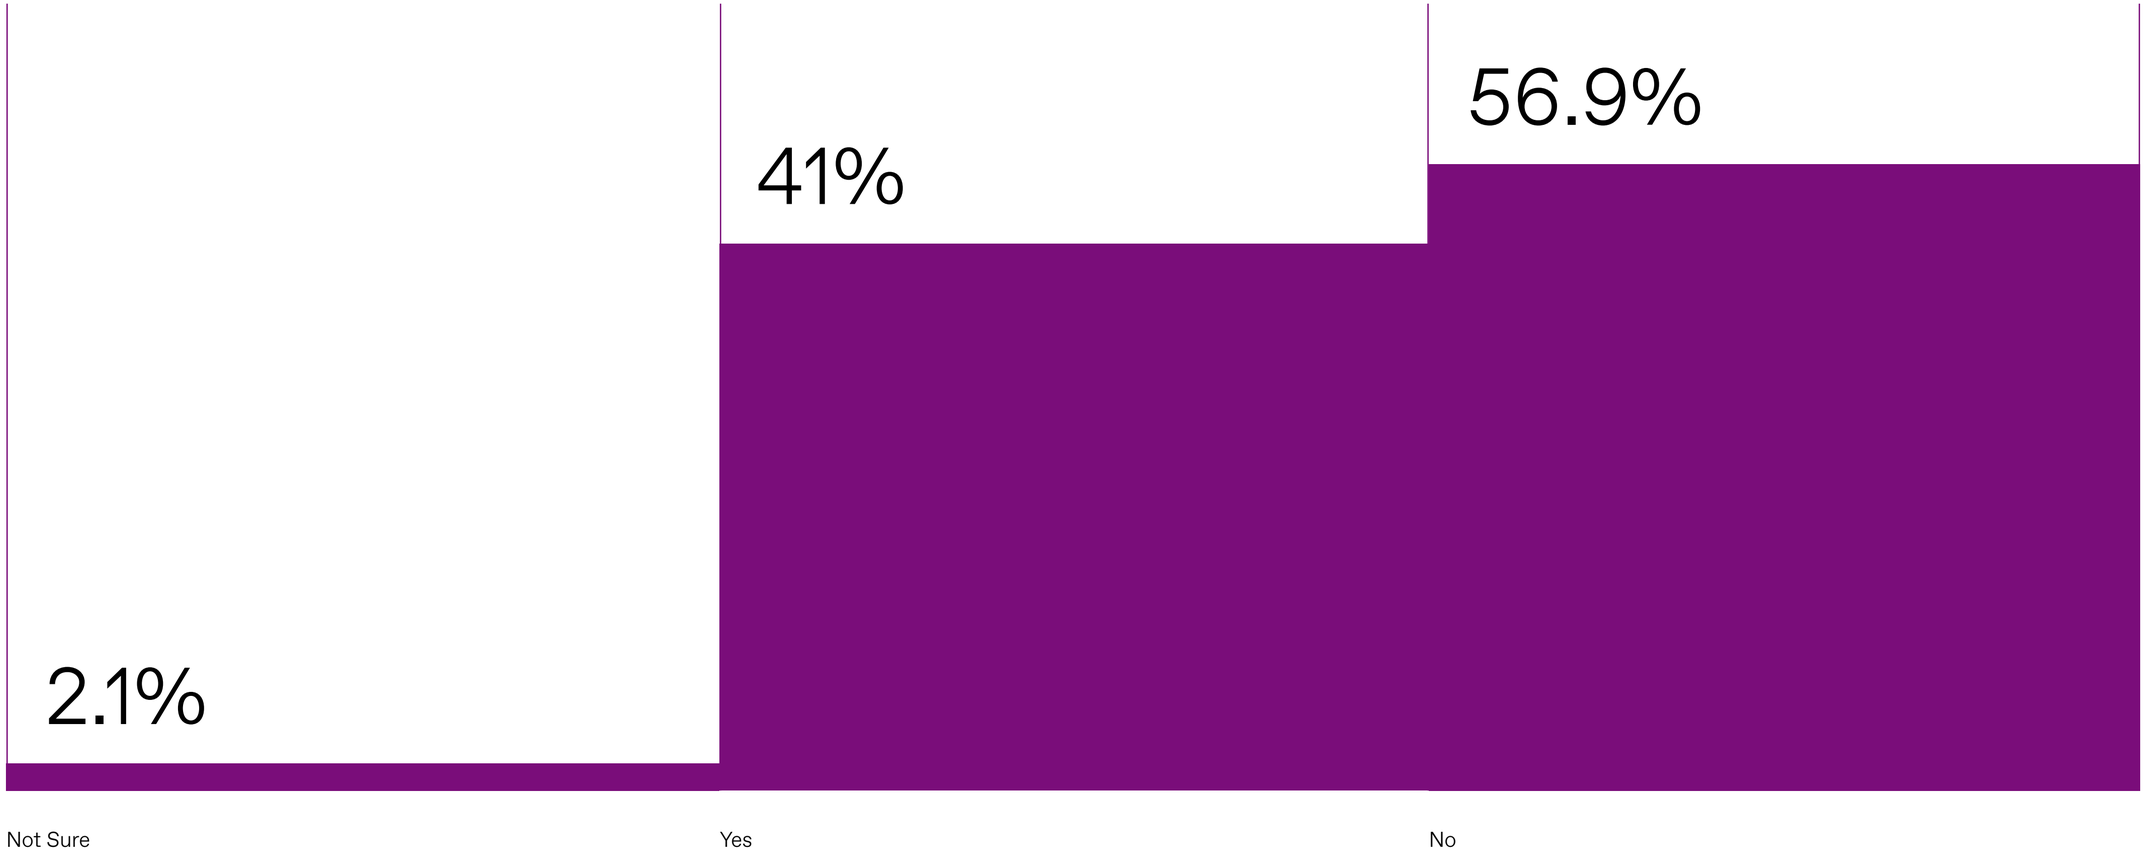


[3] https://www.axios.com/2025/08/19/north-korea-it-worker-fraud-fortune-500

# 1.5 A Stated Priority but Inconsistent Reality

**At your organization, what kind of a priority is protection against deepfake and identity fraud threats?**

Three out of four respondents said protecting against deepfake and identity fraud is already a priority with the rest stating it will be a priority within 12 to 24 months. This suggests an encouraging sense of urgency among enterprises. But responses covered later in this report suggest a lack of deep understanding of the threat or concrete, prioritized actions needed to bolster defenses.

The disconnect between awareness and true preparedness becomes even more apparent in the next section which examines how enterprises are currently protecting remote communication channels, where deepfake-enabled impersonation is becoming more prevalent and difficult to detect.

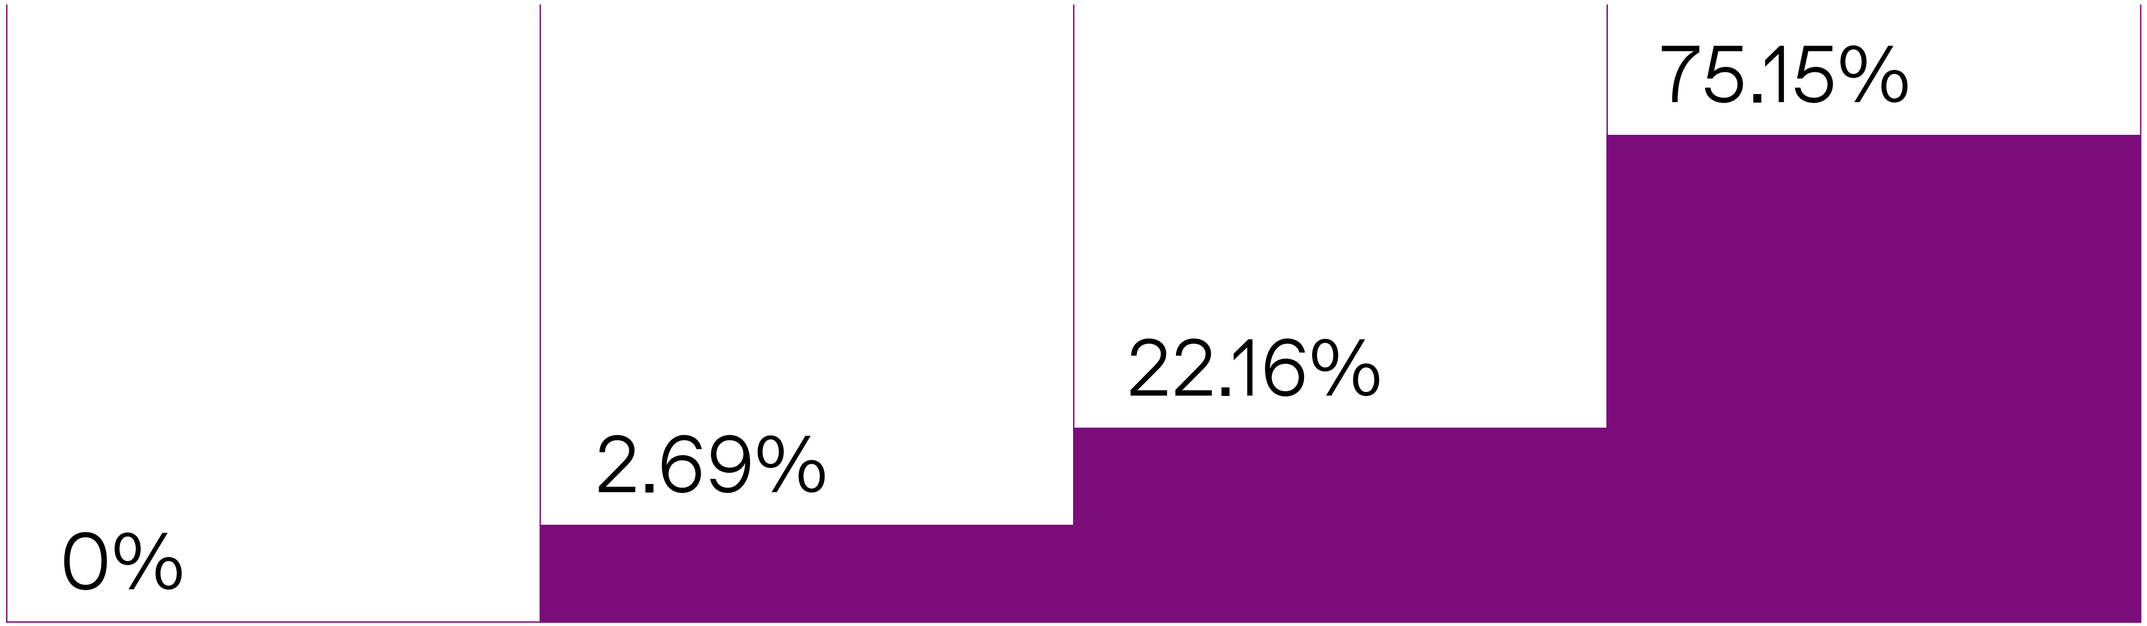

# 1.6 Response Plans in Place But Not Always Effective

**If your organization has had a recent deepfake incident, were you ready with an incident response plan?**

Only 3% of respondents admitted to not having a response plan for deepfake incidents, and two thirds of organizations that had one in place when they experienced a deepfake stated it was sufficient.

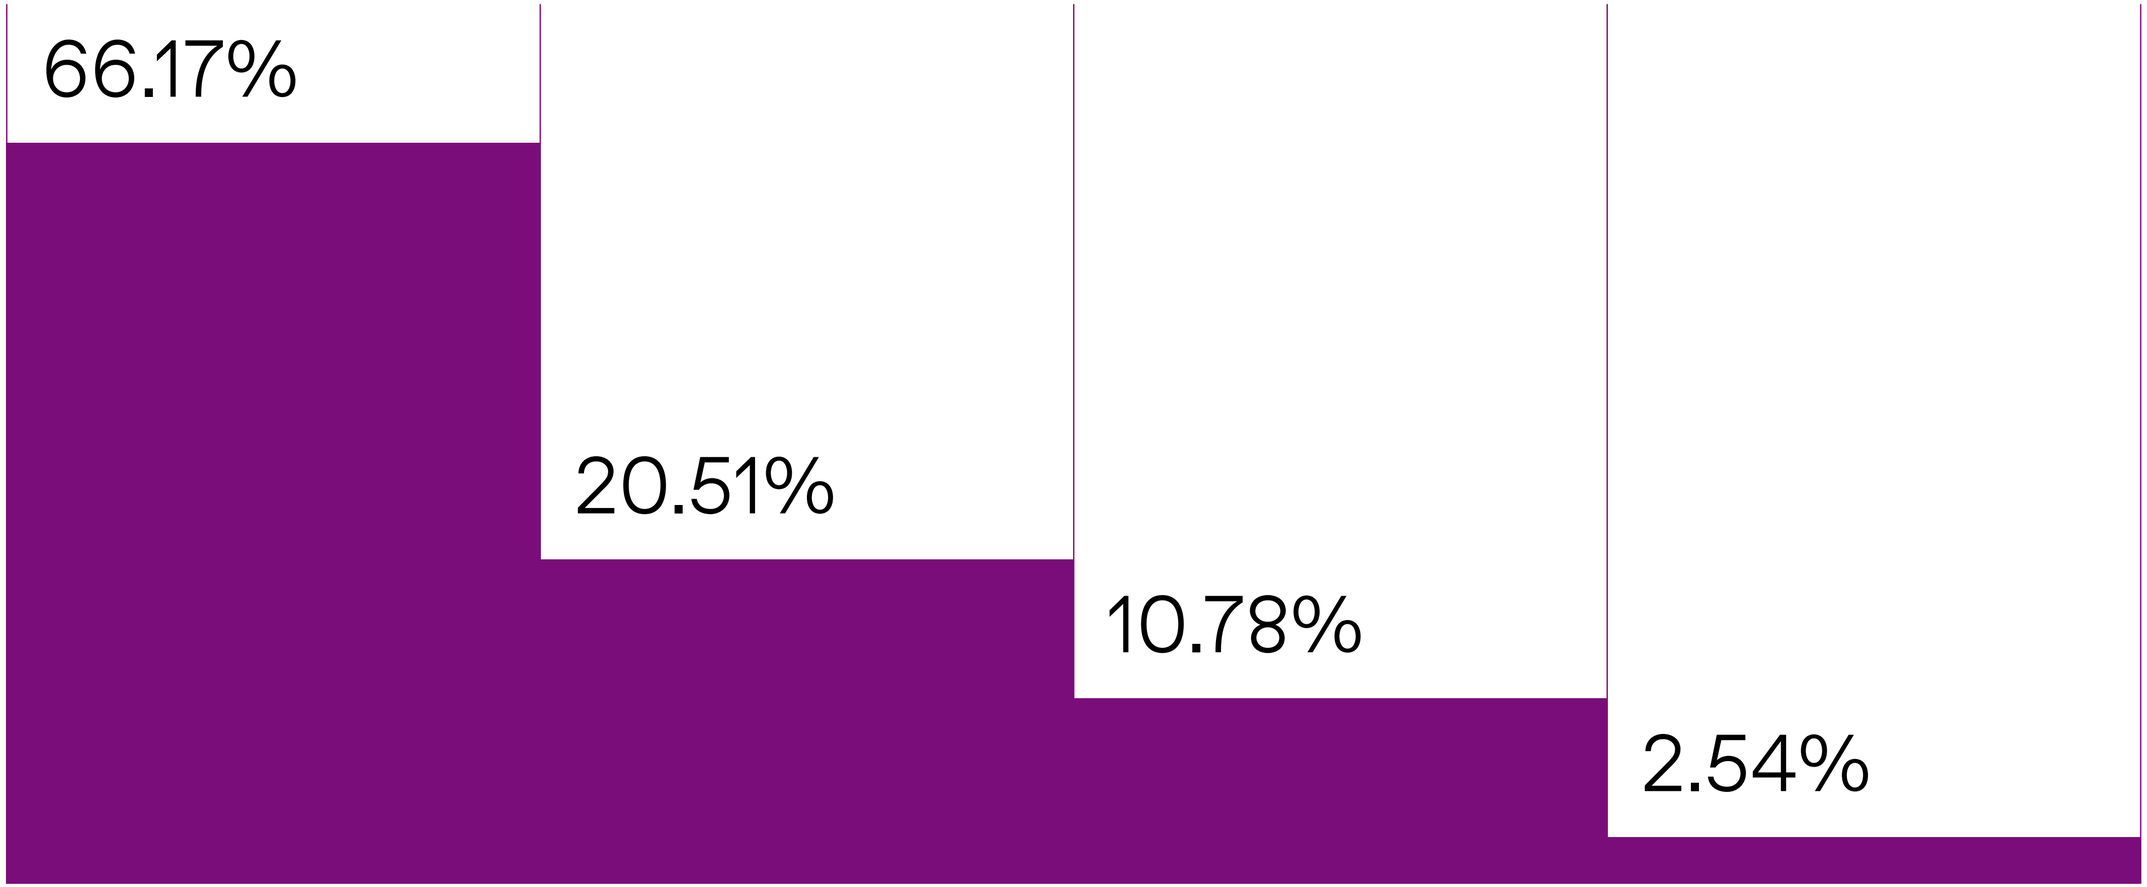

# 1.6

Employee training was the most common component of a deepfake response plan (67% of respondents), followed closely by deploying AI-based detection software (65%). Training is only one aspect of a cybersecurity program, and it can quickly become outdated as deepfake tools improve the realism of their output on a monthly, if not weekly, basis.

Rounding out the top three responses was the implementation of identification verification, though only about half of respondents said they had adopted tools to verify identity as part of their defenses.

Many identity verification tools function as a point-in-time check. They can also be bypassed with photos of identity documents and corresponding “selfie” images sold on cybercrime marketplaces or stolen from victims. [4]

Additionally, these one-time checks are rarely tied to the person who later appears on a call or in an interview (e.g., stand-in interviewees), leaving a gap between the identity that was verified and the individual that ultimately shows up on a call.

[4] https://www.forbes.com/sites/daveywinder/2024/12/27/dark-web-face-id-farm-warning-as-hackers-build-identity-fraud-database/

# 1.6

**What actions did your response plan include?**

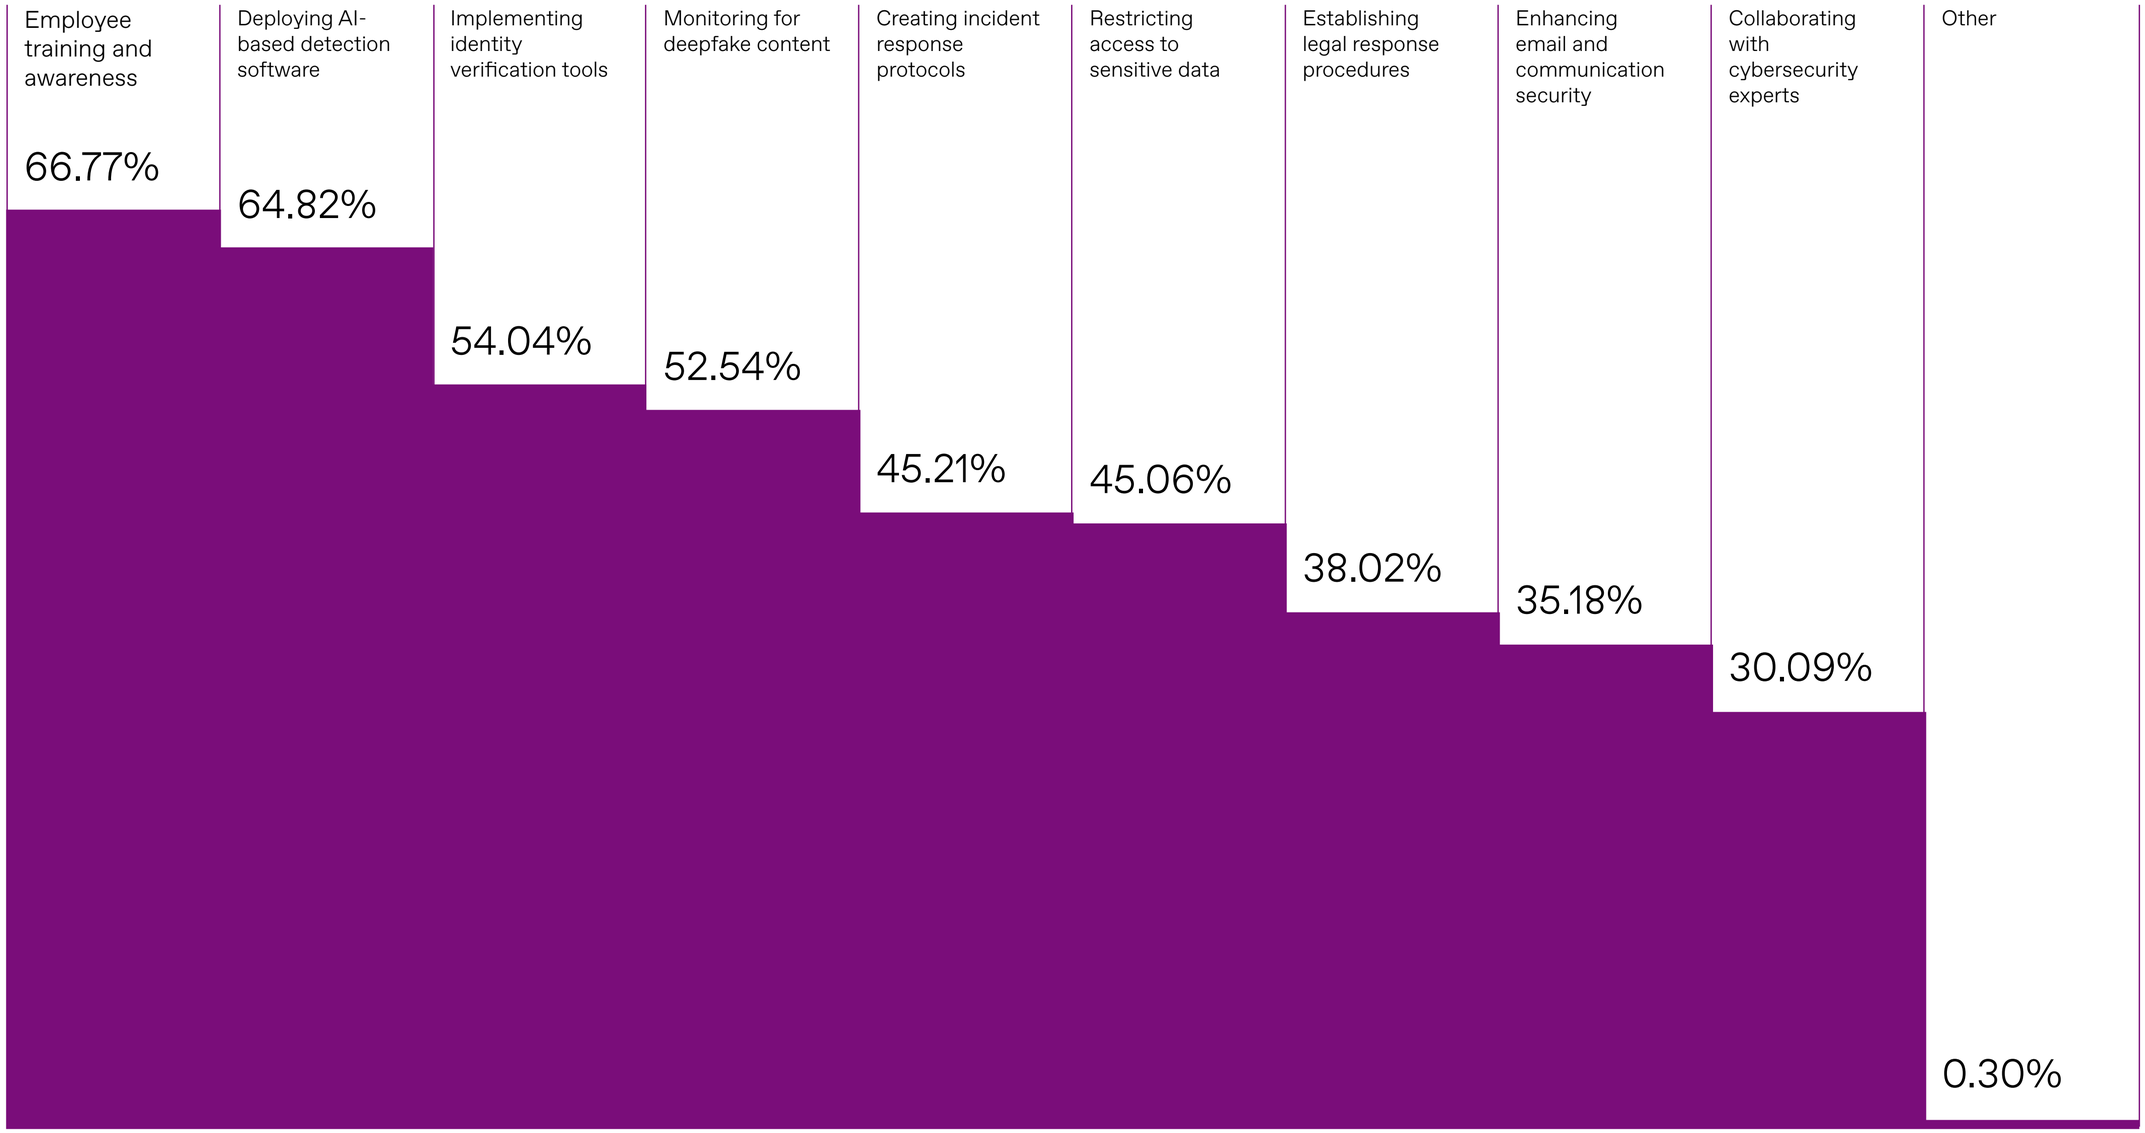

# 1.6



The data shows deepfakes targeting enterprises are a pervasive threat.

The next question is whether enterprise workflows and remote communication channels are prepared for them.

# 2.0/ Remote Communication Channels are Vulnerable

The gap between perceived preparedness and actual risk becomes particularly evident in the remote communication channels enterprises depend on every day. As distributed work becomes standard, enterprises rely on video and voice communication – channels which many times lack true identity assurance.

Remote work and collaboration has conditioned people to trust these interactions, assuming that seeing a face or hearing a voice is sufficient in verifying the identity of a person on the other end.

# 2.1 Most Organizations Use Voice and Video Calls to Verify a Person's Identity

**Does your organization use voice or video calls to verify a person's identity as a part of any business process?**

The majority (92%) of respondents said their organizations definitely, very likely, or probably use voice or video calls to verify a person's identity as part of a business process.

Real-world incidents demonstrate that fraudulent applicants are already exploiting these channels to infiltrate hiring processes, and overreliance on verbal or visual cues leaves enterprises exposed.

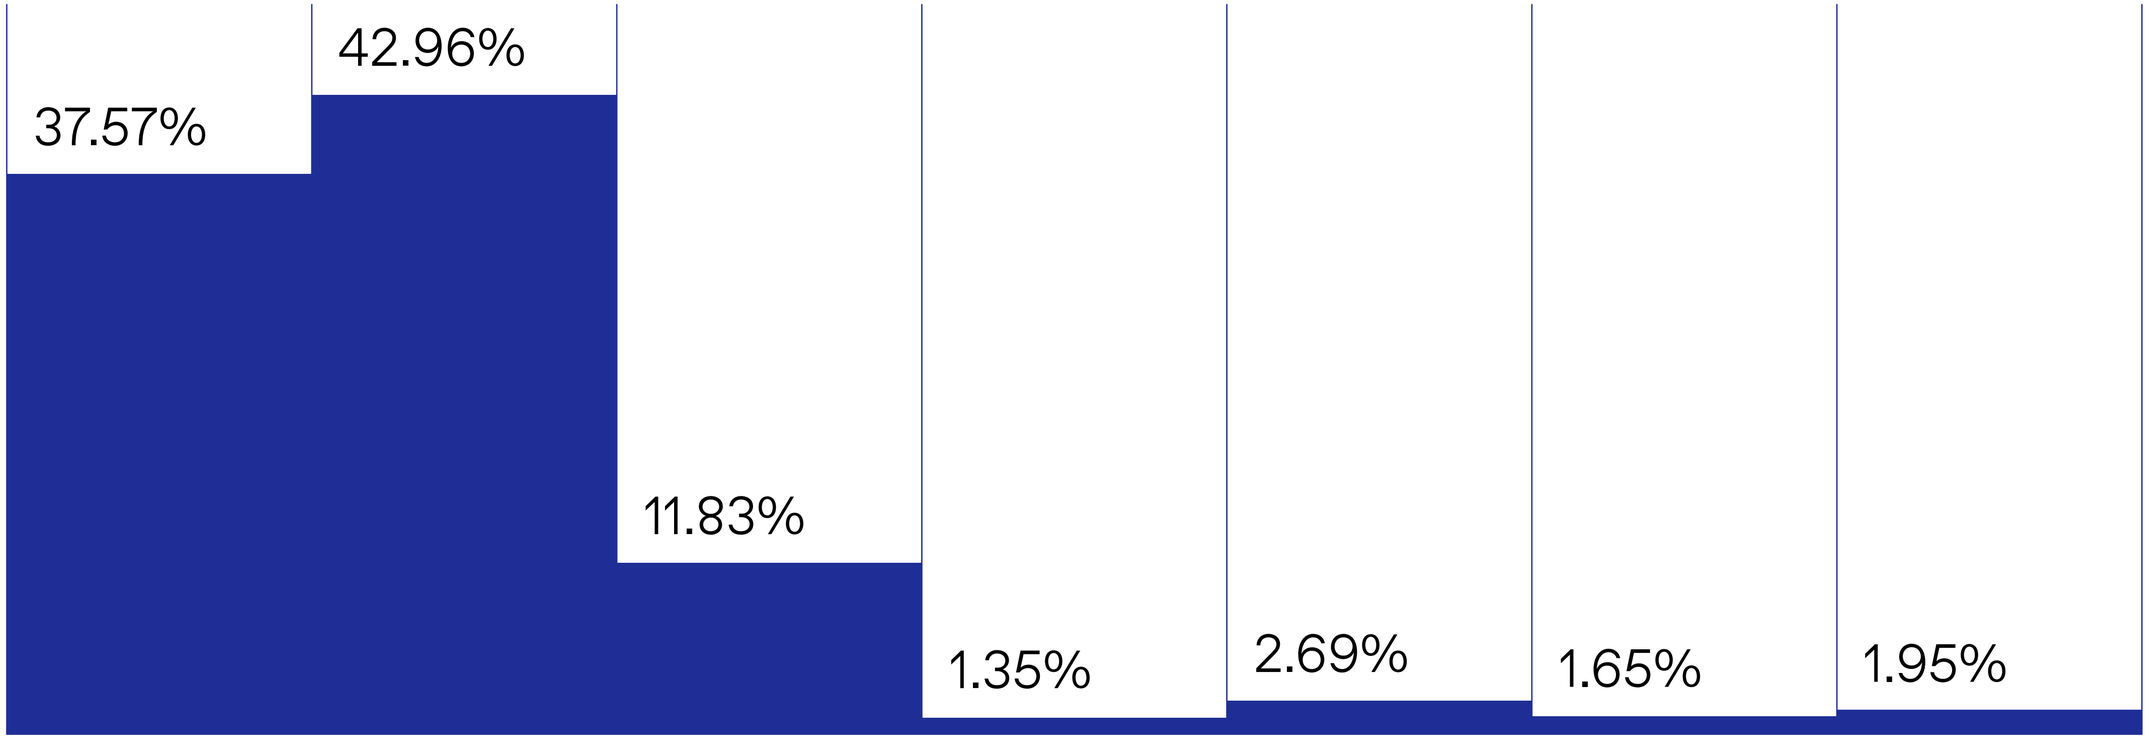

# 2.2 But These Channels Remain Vulnerable to Impersonation

Added to this, only about half of respondents always verify the identity of a person speaking during a video call (55%) or require authentication to stop impersonation attempts on video conference calls (52%). In practice, some calls will inevitably include unauthenticated participants.

The resulting gap presents adversaries with an opportunity to insert themselves into day-to-day operations, highlighting why multiple layers of defense are necessary to restore the integrity of these channels and harden them against impersonation.

**Does your organization verify the identity of the person speaking during a video call?**

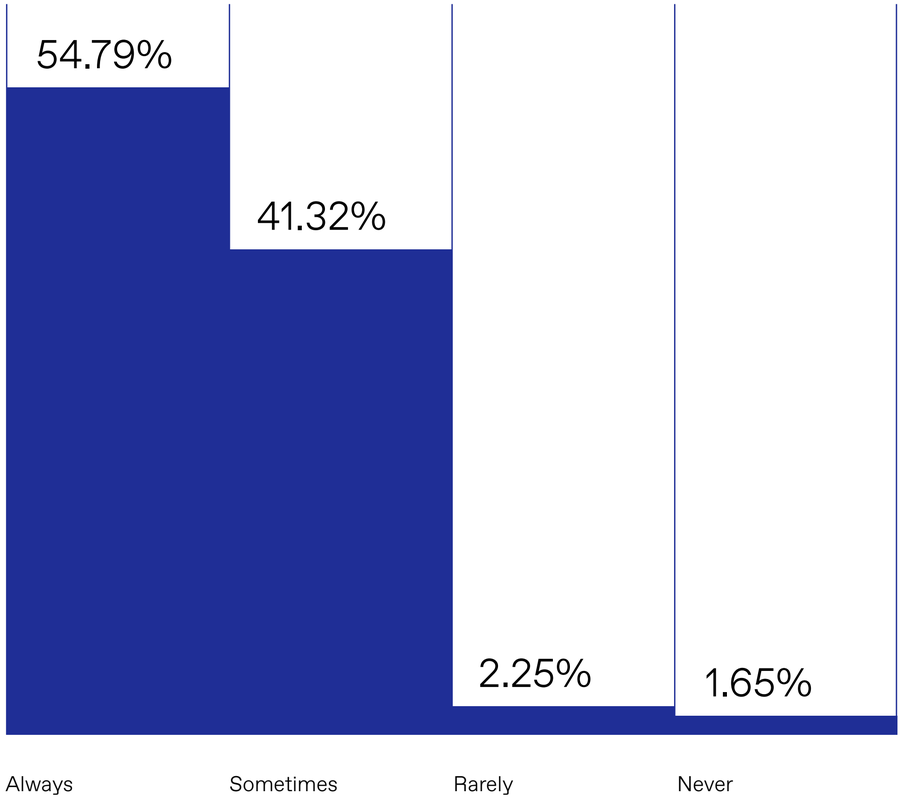


**Does your organization require authentication to stop impersonation attempts on video conference calls?**

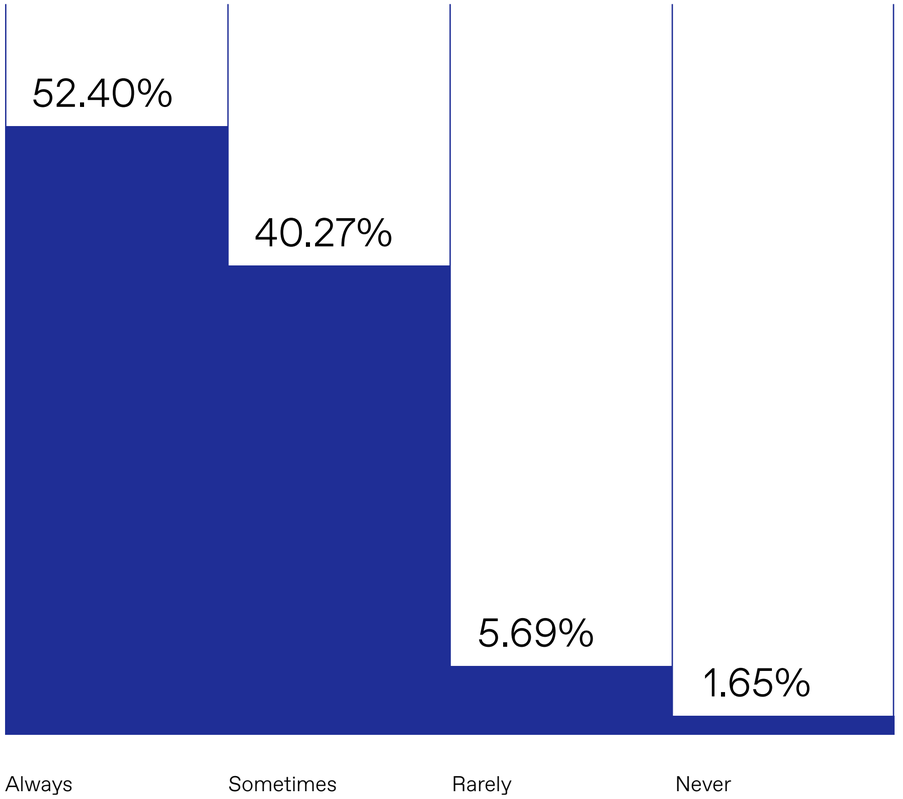

# 2.3 Confidence in Distinguishing Synthetic from Authentic Digital Media

**Do you feel confident distinguishing a real voice from a deepfake voice or face in a video conference today?**

When asked about their confidence in distinguishing a deepfake, only 2% of respondents questioned their ability to tell a real voice or face from a synthetic one. Interestingly, respondents whose organizations experience deepfakes more frequently were more likely to be confident in their ability to distinguish them.

In one academic study, researchers found that people can't reliably tell when a voice is AI-generated. Study participants were fooled by an AI-generated version of a real person's voice 80% of the time and identified a synthetic voice only 60% of the time.[5] These detection rates fall far below thresholds for sufficient enterprise protection.

This overconfidence is especially concerning in light of the inconsistent verification and authentication practices reported over voice calls and videoconferencing, which are channels attackers target to bypass technical safeguards. Effective defense against remote impersonation must go beyond human intuition to include layered, automated, continuous validation of the other parties on a call before they're trusted.

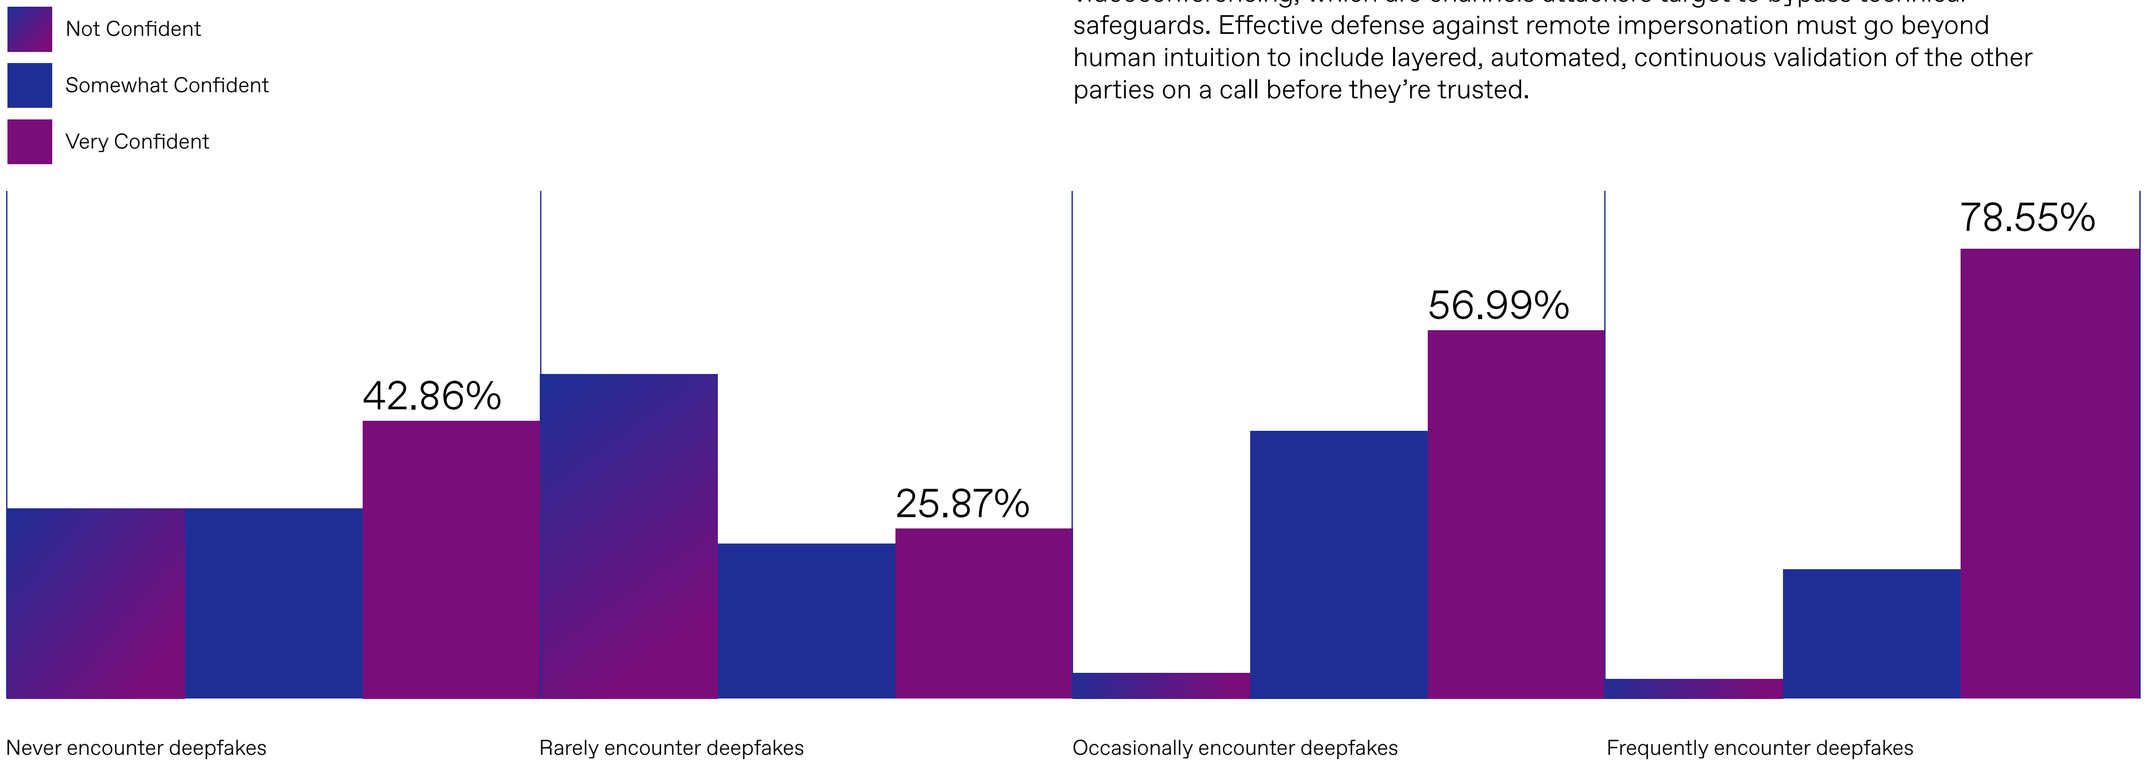


[5] https://www.nature.com/articles/s41598-025-94170-3

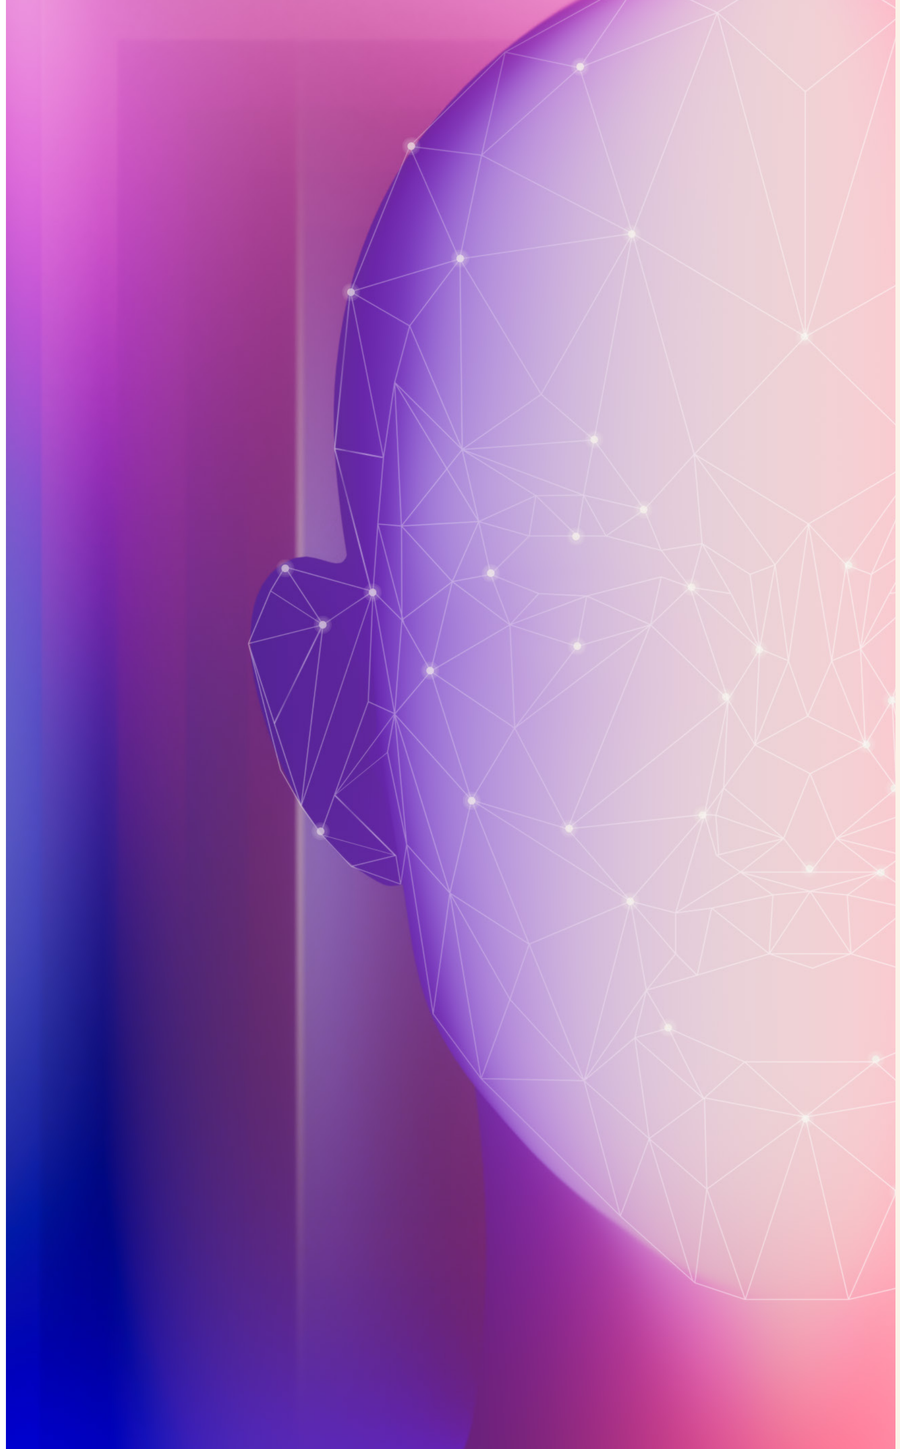

# 3.0/ Deepfakes are an Identity Problem & IAM Must Evolve

Employee credentials are often tied to an identity at a single point in time and rarely reassessed for ongoing trustworthiness. This creates a gap between the person an organization originally verified and the entity that later appears on a call. Adversaries exploit this weakness as illustrated by FBI reports of a Scattered Spider tactic in which attackers pose as employees to socially engineer corporate IT help desks. The actors seek to convince help desk staff to reset passwords or transfer MFA to attacker-controlled devices, granting them access to employee accounts and corporate systems.[6] This erosion of trust, and its static nature, underscores why identity and access management must evolve to address AI-driven deception.

[6] https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-320a

# 3.1 Overconfidence in Legacy Controls

**Do you believe multi-factor authentication alone is sufficient to stop AI-driven impersonation threats?**

Overall, 82% of respondents said their organization very likely (42%) or definitely (40%) has adequate protections in places against deepfakes and synthetic identity fraud. At the same time, many respondents place trust in multi-factor authentication as a primary defense against AI-powered impersonation attacks. Credential abuse, however, remains one of the top three initial access vectors in data breaches.[7] As the previous section illustrated, enterprises are not consistently applying authentication best practices to remote communication channels anyway, which contributes to attackers targeting these channels to sidestep authentication controls.

More than three-fourths (79%) of respondents said they definitely, very likely, or probably believe multi-factor authentication alone is enough to stop AI-driven impersonation attacks.

Misplaced confidence in MFA reflects a view of identity as static in a world where attackers can combine stolen credentials with AI-cloned voices and video to pass as legitimate users.

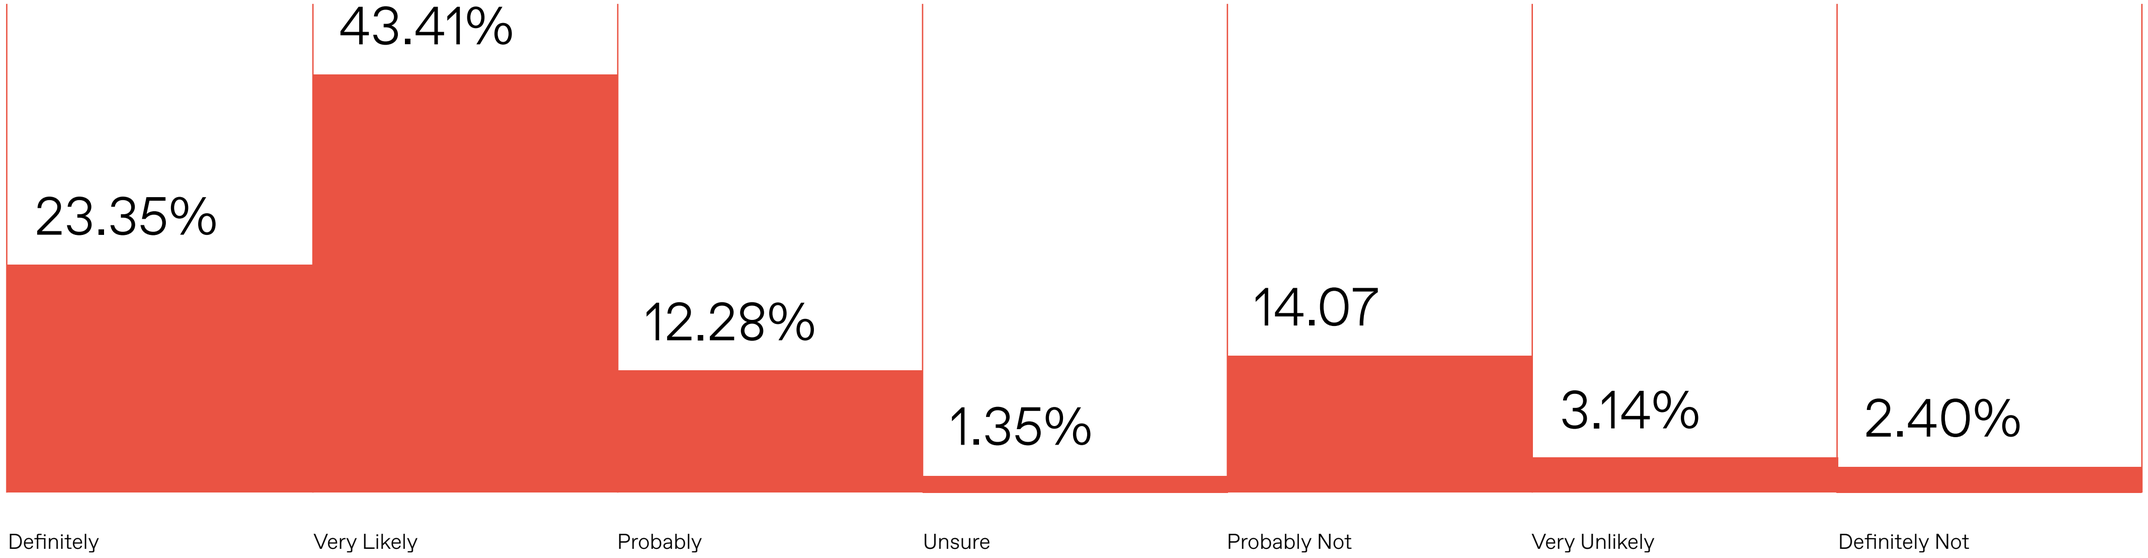


[7] https://www.verizon.com/business/resources/reports/dbir/

# 3.2 IAM Strategies Have Not Yet Caught Up to AI-driven Threats

**Have GenAI-enhanced synthetic identity threats prompted your organization to rethink its IAM (Identity and Access Management) roadmap?**

Only about half of respondents are definitely rethinking their IAM roadmaps to account for Gen AI-enhanced identity threats. Possession of valid credentials doesn't prove identity, and as established earlier, research shows human intuition alone cannot reliably detect when something "sounds off" or "looks off" during a call.

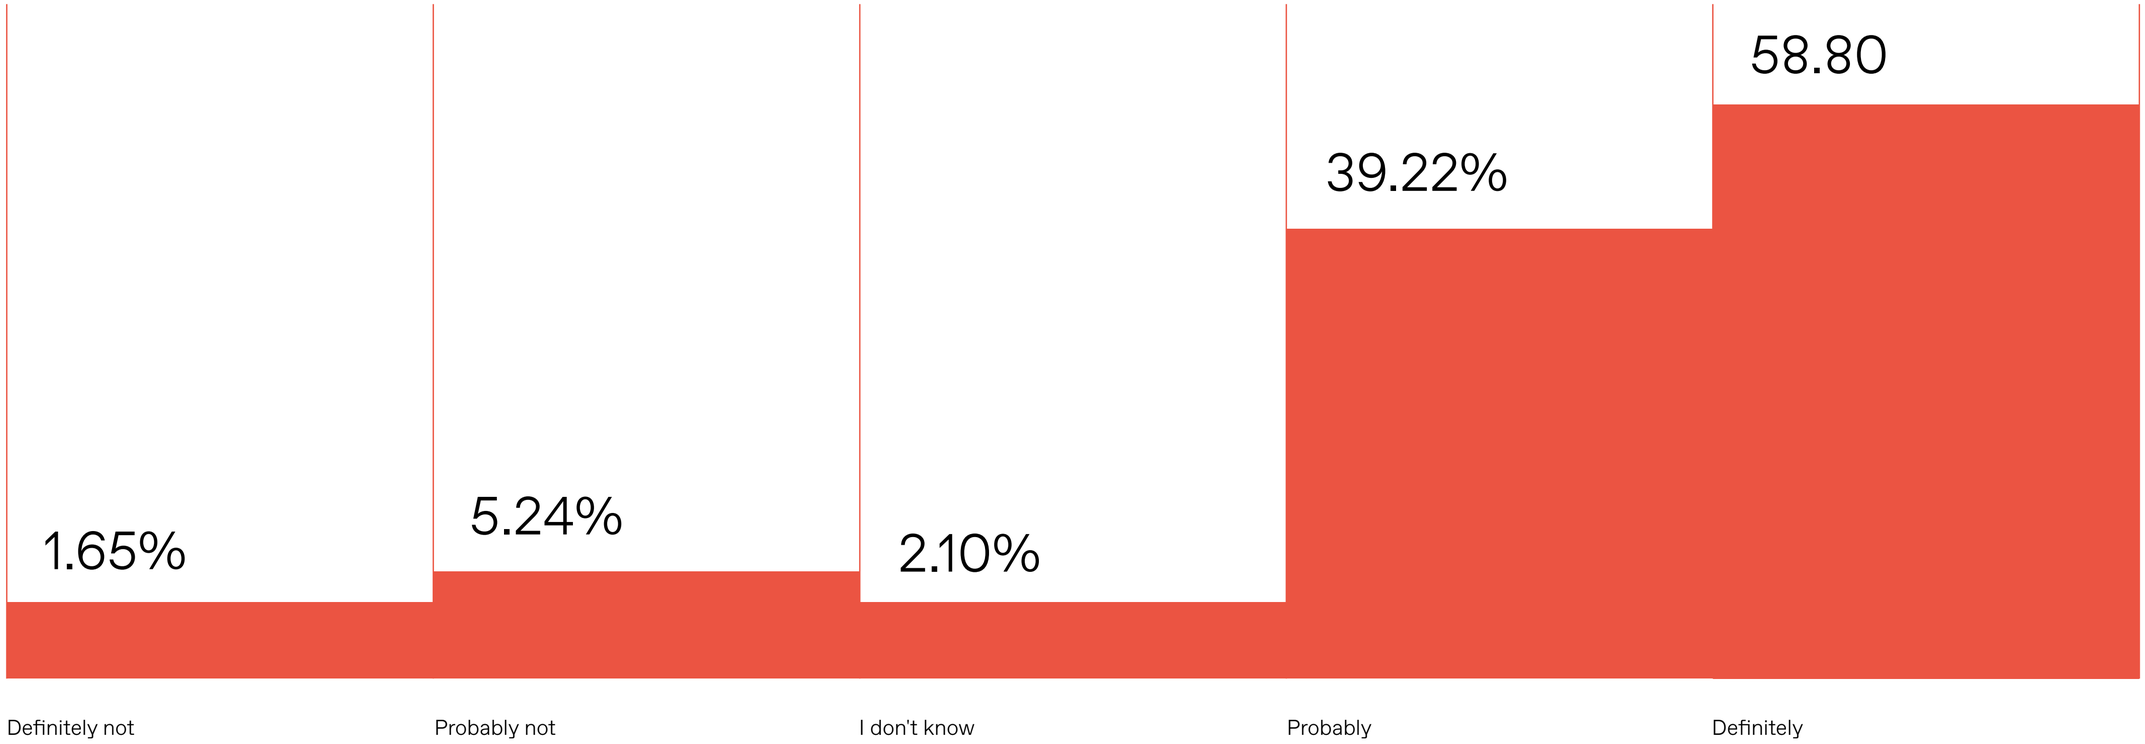

# 3.3 No Consensus in Modernizing IAM Amid Accelerating Deepfake Threat

**Which changes would be part of your organization's IAM strategy to address deepfake threats? (Select all that apply)**

No single aspect of IAM modernization reached a clear consensus among respondents. A mere 28% of respondents cited deepfake-resistant verification tools as a priority. While stronger biometric controls (63%) and expanded MFA (62%) top the list, their effectiveness in the face of AI-powered deception depends upon consistent, ongoing verification of biometric data and user identity across communications.

This continuity is essential to detecting anomalies and ensuring identity trust over time. Some respondents selecting real-time identity proofing and continuous behavior monitoring suggest recognition of the necessity of ongoing validation, but only about half of the surveyed population included these as part of their forward-looking IAM strategy.

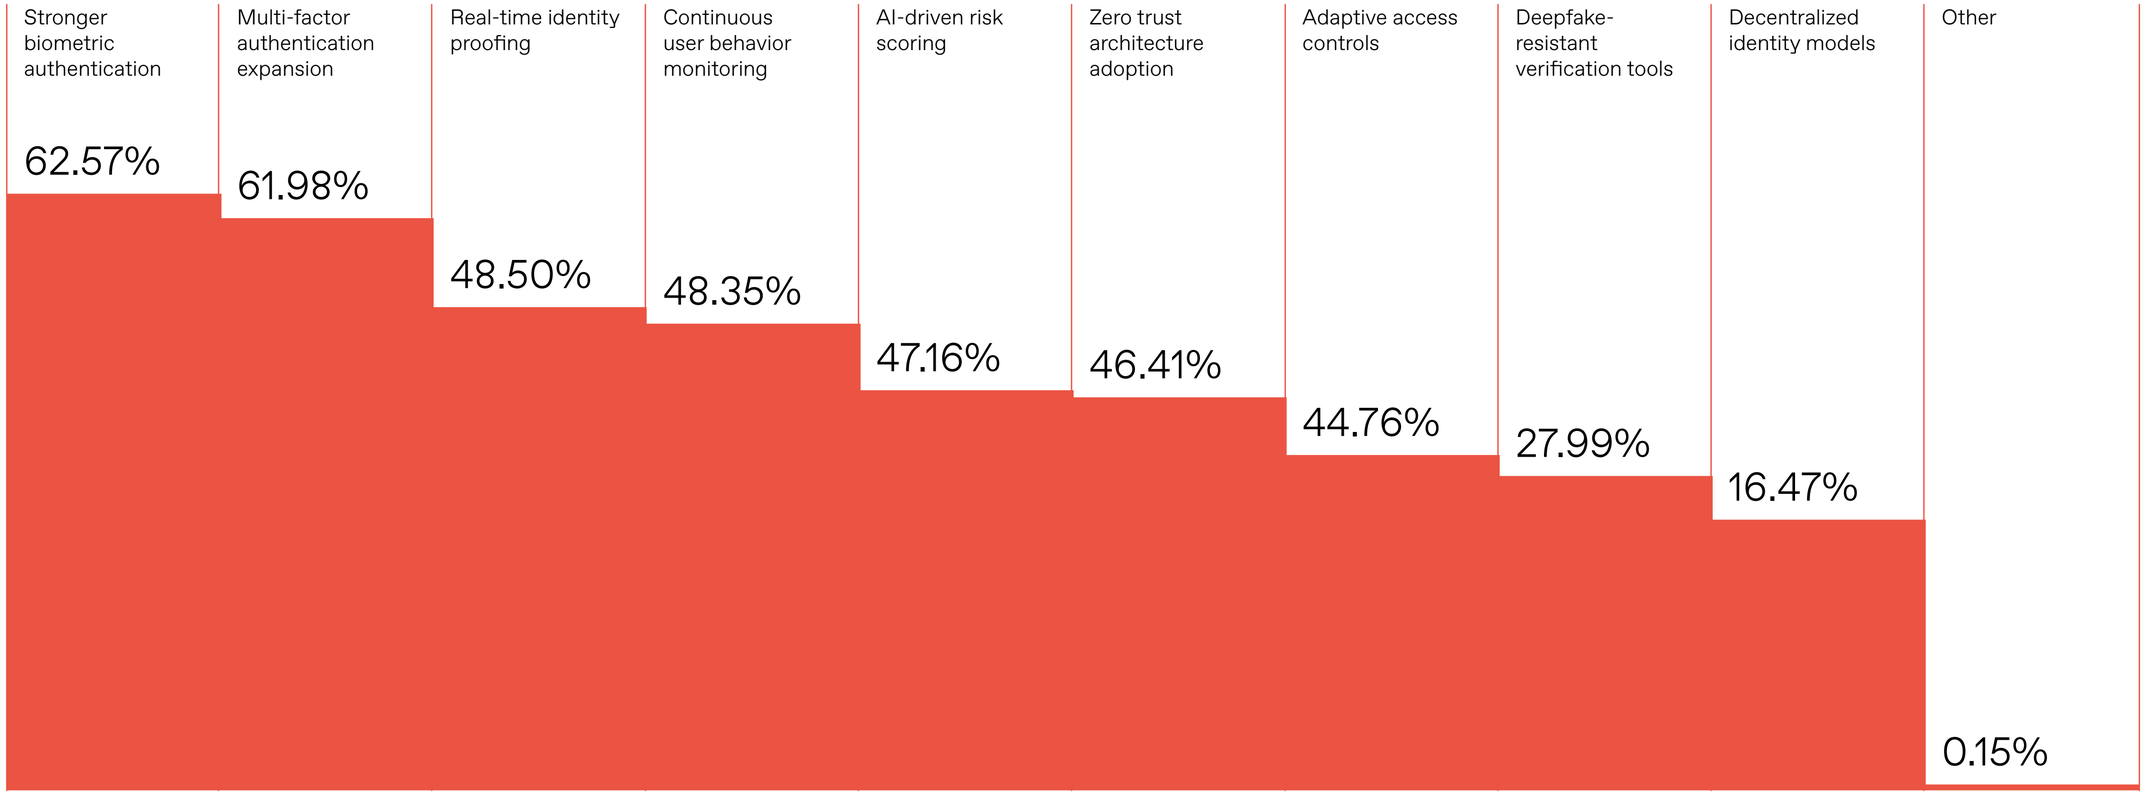

# 3.3



As deepfake tools advance faster than IAM strategies, enterprises that fail to adopt continuous monitoring and validation of remote likenesses will face a widening gap that attackers will exploit to launch impersonation attacks, carry out fraud, and gain access to corporate systems and sensitive data.

Together these findings show that while enterprises recognize identity as important to deepfake defense, current strategies treat identity as static, which leaves gaps that AI-powered impersonation can exploit.

# Conclusion



With synthetic images, prerecorded media, and real-time audio and video now easily produced by Generative AI, enterprises must rethink how they establish and maintain digital identity and trust.

Point-in-time authentication and validation rely too heavily on static checks that can’t keep pace with quickly evolving, AI-driven impersonation. With employee and customer likenesses more widely available across social media and cybercrime markets, defeating AI impersonation requires enterprises to understand their people more deeply than Gen AI can.

Because biometric authentication can now be convincingly spoofed by deepfakes, ongoing identity protection and validation become essential. A continuous identity protection approach enables enterprises to build a more robust, consistent understanding of each employee’s, partner’s, or customer’s likeness in order to detect anomalies.

# A modern approach to today's identity attack surface requires answering three core questions:

Is the presented identity on the other end of this interaction authentically human or AI-generated?

And this must be evaluated with the understanding that "AI-generated" is not always malicious, especially as agentic AI becomes more common.

Is the presented identity who or what it claims to be?

Answering this question requires ongoing monitoring of biometric and behavioral consistency across channels and over time, instead of one-time verification.

If the answers raise suspicion, what action should be taken?

Organizations need clear policies and automated enforcement mechanisms when an identity violation is surfaced.

To protect digital collaboration, enterprises must align people, processes, and technology to answer these questions quickly, consistently, and at scale across every digital interaction between their employees, partners, and customers.

---

Organizations that act now to shift from point-in-time checks to ongoing identity integrity monitoring better position themselves for resilience against AI-powered attacks.

This is why GetReal Security built its Digital Integrity Platform to provide the foundation for continuous identity protection and visibility across the enterprise and every remote interaction.

This includes validating the integrity of job applicants throughout the hiring process to ensure only authentic candidates are onboarded, protecting corporate IT help
desk and customer support staff from impersonation attempts during account recovery, and protecting
employees from executive impersonation schemes aimed
at executing unauthorized payments, stealing sensitive data, or other objectives.

# Methodology

This report is based on a commissioned, double-blind 41-question survey of 668 IT, cybersecurity, fraud, and risk leaders fielded in September 2025.

Respondents included directors (57%), vice presidents (9%), and executives (34%) at organizations with 1,000 to more than 10,000 employees spanning 15 industries with the largest representations coming from technology, financial services, pharmaceuticals, manufacturing, and healthcare.

# About GetReal Security

GetReal Security is the cybersecurity leader in detecting and mitigating threats posed by deceptive AI content including deepfakes, impersonation, and synthetic identities.

Our enterprise-class digital integrity platform combines advanced detection, forensics expertise, and threat intelligence to prevent fraud, maintain compliance, and restore trust.

It delivers unmatched visibility through dashboards revealing AI-powered manipulation, the people impacted, and conversations affected, all seamlessly integrated with existing enterprise technology stacks.

GetReal