IRONSCALES

IRONSCALES Fall 2025 Threat Report

Beyond Detection: The $280K Reality of Deepfake Attacks

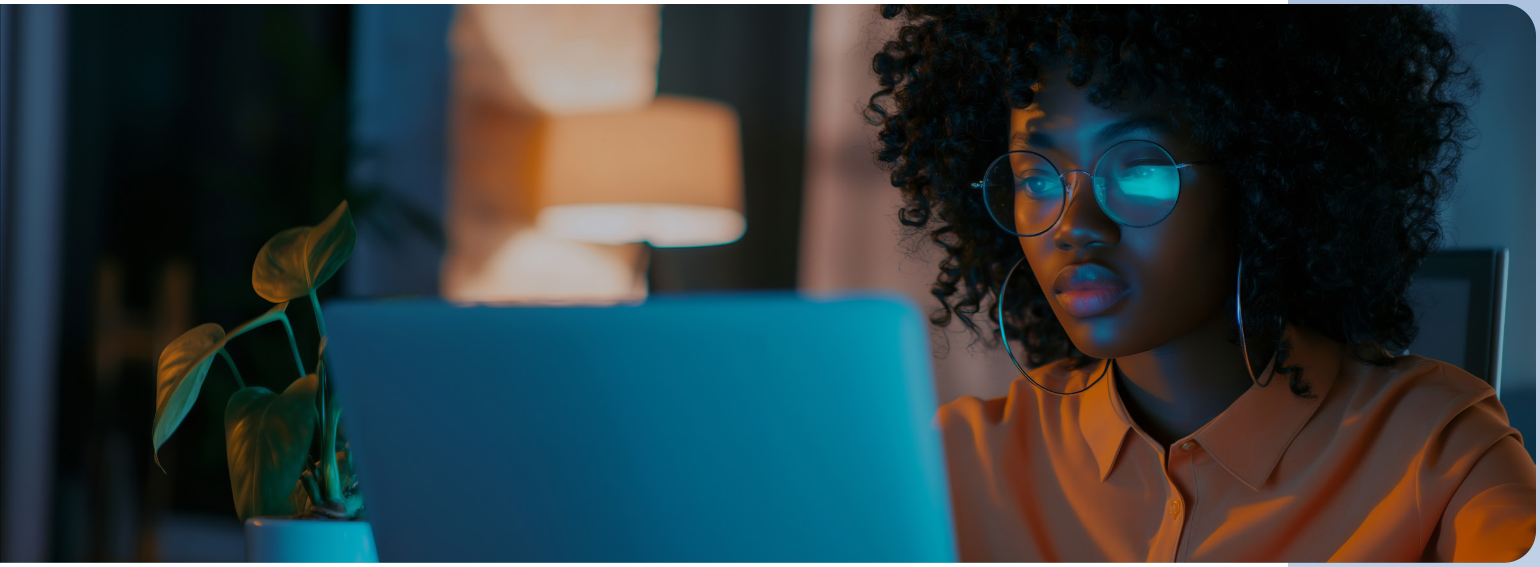

Executive Summary & Key Takeaways

In our second annual Fall Threat Report, produced with market research firm Censuswide, one truth rises above the noise: **the era of deepfakes and deepfake-related incidents is here.**

Deepfake-related[1] incidents increased 10% year-over-year, **with 85%** of surveyed IT and cybersecurity professionals reporting their organizations experienced one or more incidents in the past 12 months. Over 40% faced three or more attacks, and 55% of respondents reported suffering **financial losses averaging more than $280,000.**

Deepfake-related incidents are here and making their presence felt across industries. With the underlying AI models being used to generate deepfakes improving at breakneck speed, these attacks have only grown increasingly sophisticated. These varied attacks are rampant, costly, and leading to mounting concern among respondents.

The data reveals that there is a rift between perceived defensive capabilities and actual preparedness for the rising tide of deepfake-related incidents. Respondents are stepping up targeted training programs, but the results have thus far been lackluster. At the same time, investments (both current and planned) in deepfake defense are seeing an uptick, but still there appears to be a lack of urgency among respondents.

In this report, we will share the data that illustrates this new reality, and offer recommendations on what organizations can do to adapt. The following **five key takeaways** offer a snapshot of the report's findings and what they mean for security teams moving forward:

[1] Deepfake-related incidents encompass attacks using synthetic or manipulated media, from basic image alterations and doctored emails to sophisticated real-time voice and video impersonations.

Key Takeaways

1 Attacks Are Widespread, Growing, and Costly

Deepfake attacks have arrived. They are now being encountered by the majority of respondents regardless of industry, and have serious financial consequences. In our survey, we found that:

- **85% of respondents report having experienced one or more deepfake-related incidents within the past 12 months**, with over 40% experiencing three or more attacks.
- **55% of respondents reported losses attributed to** deepfake or AI-voice fraud in the past 12 months.
  - Over **61% of organizations** that have lost money in a deepfake attack report losses in excess of $100,000, with nearly a fifth **(19%) reporting having lost half-a-million or more.**
  - Mean losses stand at over $280,000 for deepfake-related incidents.
  - Over **5% of organizations have lost $1 million dollars** or more to deepfake-related incidents.

2 Deepfake Defense Is Rapidly Climbing the Ranks of Cybersecurity Priorities, But Investment Lacks Urgency

In light of the above findings, it should come as little surprise that deepfake defense is rapidly becoming a top priority for security teams. However, we are seeing a significant lag in actual investment. While both current and anticipated investment in deepfake defense is increasing, the gap between concern and action reflects market dynamics: deepfake defense solutions are relatively new to the market, and many organizations are waiting for more mature, proven technologies before committing budget, even as losses mount.

- Nearly all (99%) of respondents say that deepfake defense will be important to their cybersecurity strategies over the next 12-18 months.
- An overwhelming **71% of respondents say deepfake defense will be a top priority** during that time (up from just 43% in 2024).
- And yet, **only 37% of respondents said their organizations are already investing in deepfake defense**.
- A mere 1% said they had no plans to invest, down from nearly 17% in 2024.

3 While Email and Static Images Stand Out, Other Vectors Are Close Behind and Quickly Catching Up

Static images and email-based attacks remain the most prevalent modalities for deepfake threats. However, other more sophisticated forms are quickly catching up, with huge leaps in volume occurring in the past year alone.

- Email-based deepfake attacks are tied with static image manipulation as the most common threat vector at 59%[2] each.
- However, other vectors are making up ground quickly:
    - Recorded content: **audio/voice manipulations rose from 25% to 52%; videos from 33% to 45%**.
    - Real-time attacks: live video manipulation increased from **30% to 41%**; live voice-only calls showed identical growth.

4 Training Is Widespread but Its Efficacy Is Questionable

More organizations are investing in deepfake-specific cybersecurity training and testing. However, the efficacy of these training programs remain questionable, as both real-world success rates and simulation testing results remain disappointing.

- **88% of organizations have provided deepfake-related cybersecurity training** (up from just 68% in 2024).
- 11.6% have provided no training.
- Only **8.4% of organizations saw first-try simulation pass rates of 80% or above.**
- The first-try deepfake-simulation pass rate average was 44%, suggesting a significant skills/retention gap worth addressing.

[2] Attack vectors referenced include: static image manipulation (doctored photos, altered documents, fake profile pictures); email-based attacks (phishing messages containing any form of deepfake content); recorded audio/voice (pre-created voice messages, voicemails, or audio files); recorded videos (pre-made video content with manipulated faces or actions); live videos (real-time face/appearance manipulation during video calls); and live audio/voice (real-time voice cloning or alteration during phone calls).

5 The Preparedness Paradox: While Concern Deepens, Defensive Overconfidence Grows

This year's survey shows near-universal concern about deepfake-related threats, with 94% of IT professionals expressing at least some level of concern. Yet a paradox persists - despite widespread incidents and reported financial losses, an overwhelming 99% of respondents claim confidence in their defenses.

- **94% of respondents expressed at least some level of concern** about the threat deepfakes currently pose to their organizations.
- Over **63% reported being "very concerned"** about the threat they pose (a 15+% increase from last year).
- Nonetheless, an overwhelming **99% claim to be confident** in their defenses.
- And the degree of confidence is growing: The percentage of respondents saying they were "very confident" has shot up by over 30% in just one year, from just over 40% in 2024 all the way to 73% in 2025.

Main Report

Section 1

Deepfake-Driven Attacks Are Here and Taking a Heavy Toll

Deepfakes are no longer a novel or emerging threat. They are increasingly ubiquitous and taking very real financial tolls on organizations. In this year's survey, the overwhelming majority (85%) of professionals polled reported their organizations experiencing one or more deepfake-related incidents within the past 12 months—a 10% increase compared to just one year prior. What's more, well over 40% of respondents reported experiencing three or more attacks in the same time frame. This represents a sizable increase compared to just one year ago, and illustrates the dual trend of increasing attack frequency and volumes.

Overwhelming majority

85%

of professionals polled reported their organizations experiencing one or more deepfake-related incidents wihin the past 12 months.

Q2. In the last year, how often has your organization encountered deepfake-related incidents?

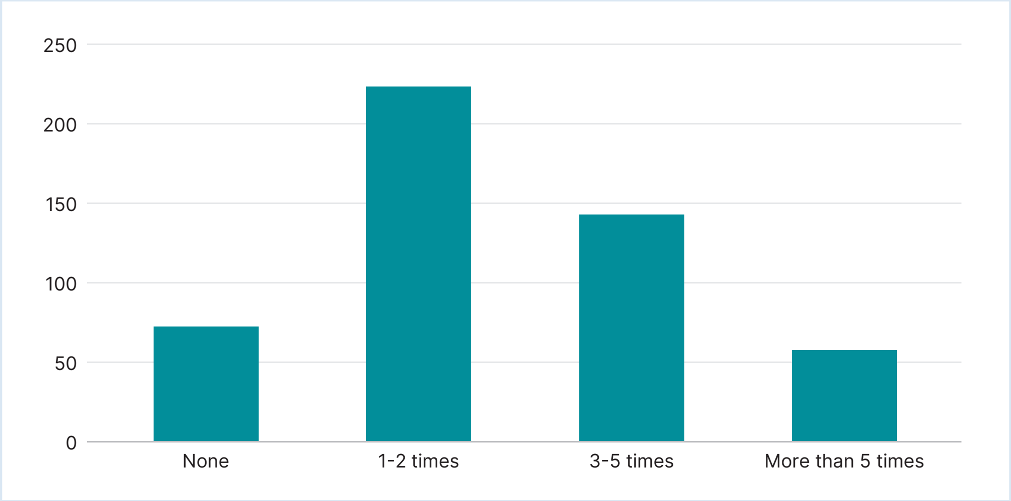

Worse yet, however, is the rate at which these attacks are proving successful. WIth well over half (55%) of the organizations polled reporting financial losses as a result of deepfake or AI-voice fraud in the past 12 months, respondents are still unprepared for this rising threat. And these financial losses are significant to say the least—with six-figure losses becoming normative (and many losses likely going underreported). Over 61% of organizations that have reported losses from deepfake-related incidents in excess of $100,000, with nearly a fifth (19%) reporting having lost half-a-million or more. The average monetary loss from financially impacted organizations stood at over $280,000—no small fee for an attack vector that didn't even exist just a few years ago.

Q11a. Please specify the estimated amount of money your organization has lost to deepfake or AI voice fraud in the past 12 months?

*Those whose organization lost money to a deepfake or AI-voice fraud incident in the past 12 months

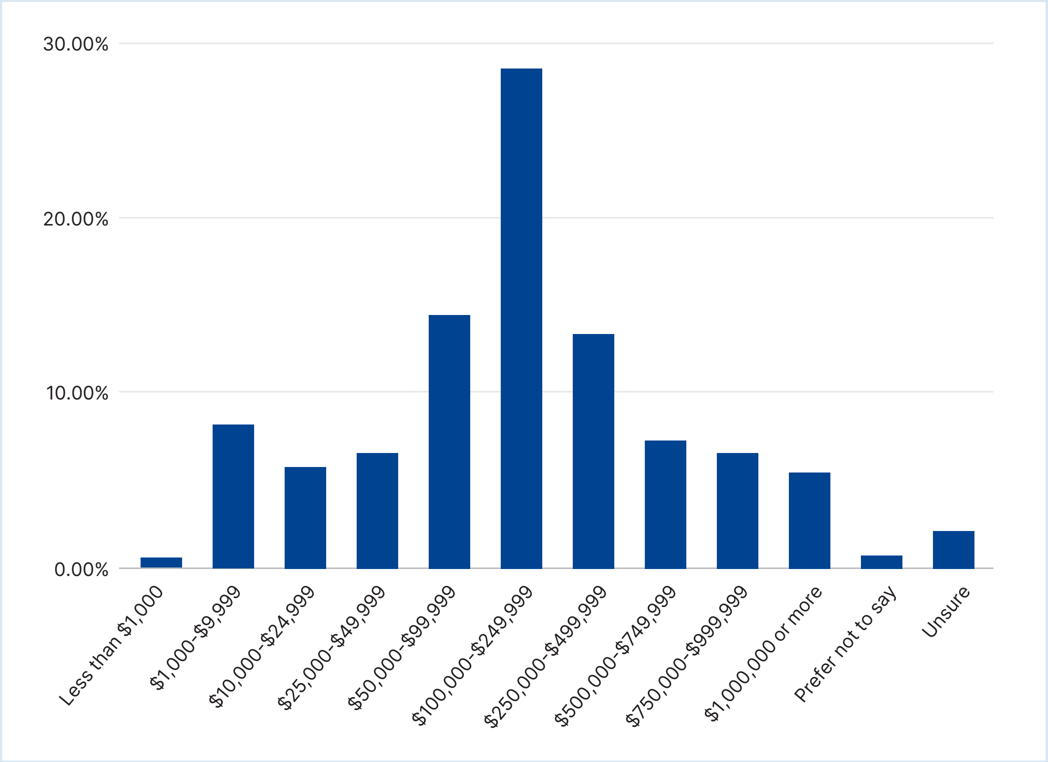

61%
of organizations that have reported losses from deepfake-related incidents in excess of $100,000.

To stem the tide, organizations must take immediate steps to harden their defenses against deepfake-related incidents. Increased investment in tailored technologies and thoughtful policies will deliver an adequate return on investment as the average cost of deepfake fraud only continues to increase.

Section 2

Deepfake Defense Climbs List of Priorities, But Investments Lag Behind

The good news is, deepfake defense is rapidly becoming a top priority for more and more organizations across industries. In this year's study, nearly all (99%) of respondents said that deepfake defense will be important to their cybersecurity strategies over the next 12-18 months, and an overwhelming 71% say it will be a "top priority" over that same time frame. This second figure represents a nearly 30% increase compared to 2024, clearly illustrating a growing unease around this dangerous new threat.

Q8. In the next 12-18 months, how does deepfake defense rank among your organization's security priorities?

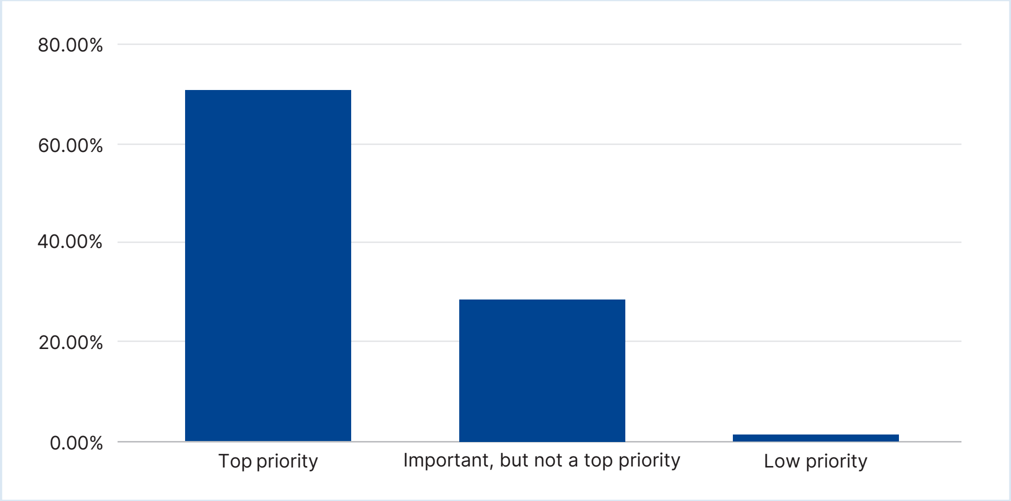

63%
of organizations still haven't invested a single dollar into deepfake defense.

With that being said, there seems to be at least some degree of disconnect between the degree of concern we're seeing in organizations and the actual investments being made to combat deepfake-related incidents. Although we are seeing increased investment (and intention to invest) overall, the overwhelming majority (63%) of organizations still haven't invested a single dollar into deepfake defense.

When measuring intent, however, a slightly more encouraging image emerges. A mere 1% of respondents reported having no intention of investing in deepfake-related defenses, a significant reduction from last year's findings, which saw 17% with zero intention of investment. At the same time, the urgency behind these intended investments leaves room for improvement. Less than a third (32%) say they plan to invest within the next five months, and an additional 22% say they plan to invest within 6-11 months. One silver lining is the fact that over 80% of respondents said their boards had requested a briefing on AI-generated deepfake risk in the past 12 months, indicating early signs of investment intent from the top-down.

Q9. When does your organization plan to invest in deepfake protection, if ever?

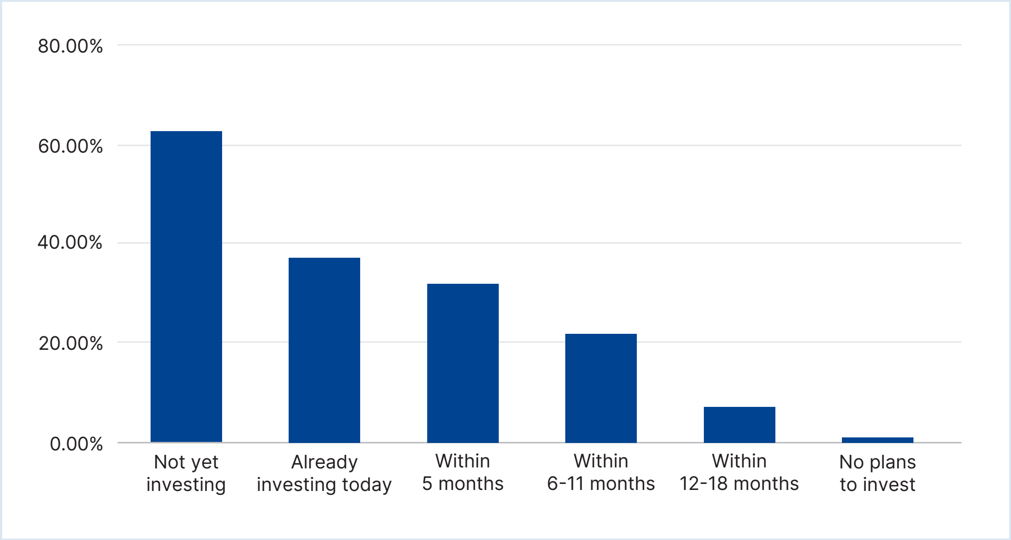

This investment gap reflects the reality of an emerging market. While deepfake defense solutions are becoming available, many organizations are understandably cautious about investing in relatively new technologies. However, with reported average losses exceeding $280,000 per incident, the cost of waiting may outweigh the risks of early adoption. Organizations should evaluate whether their current security stack can adequately detect synthetic media, and consider targeted investments aligned with their highest-risk vectors, noting that email and static images currently represent the most common attack vectors at nearly 60% each.

Given the other findings in this report, and the real-world costs being incurred by organizations, it's hard to argue that the time to invest in deepfake defenses is anything but now. Organizations would be wise to recognize this disconnect and begin investing strategically in deepfake-specific technologies, training, and policies. Perhaps, begin by looking at which departments are most at risk and allocating time/resources strategically.

80%

of respondents said their boards had requested a briefing on AI-generated deepfake risk in the past 12 months, indicationg early signs of investment intent from the top-down.

Section 3

While Email and Static Images Lead, Other Modalities Are Quickly Closing Ground

While email deepfakes and static image manipulation remain the number one most commonly reported deepfake-related modalities (both tied at 59%), other more sophisticated types of media are seeing dramatic increases—ultimately bringing them nearly on par with these more traditional threat types. In the 2024 Deepfake Threat Report, just 25% of respondents reported experiencing recorded audio/ voice manipulations. Just one year later, in the present study, that number shot all the way up to 52%, more than doubling in just one year's time.

Other modalities have seen similar upward trends, with recorded videos rising from 33% in 2024 to nearly 46% in 2025, and live video increasing by over 11% year-over-year (from 30% to 41%). Live audio and voice manipulations saw an identical increase. Collectively, all surveyed modalities saw annual increases of at least 11% or more.

25%
of respondents reported experiencing recorded audio/ voice manipulations.
In this year's survey, that number shot all the way up to
52%
more than doubling in just one year's time.

Q3. Which types of deepfakes has your organization encountered in the past year? (Select all that apply)

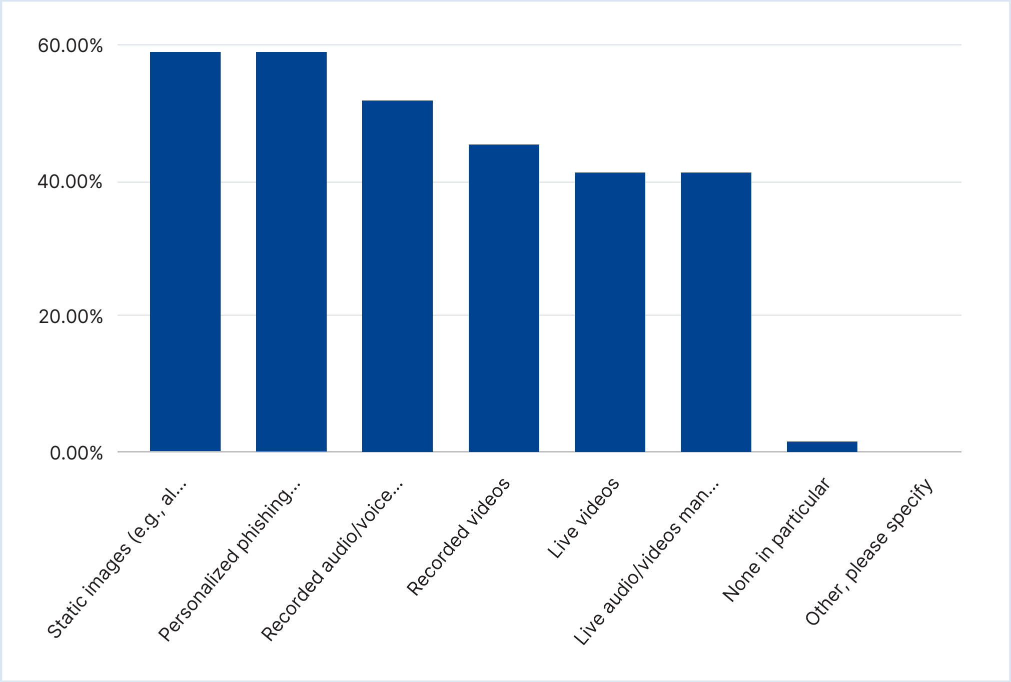

As deepfake generation tools become more sophisticated, widespread, and accessible, organizations should be cautious about over-indexing against any one single threat. While email should remain one's #1 priority, all forms of deepfakes are now real, prevalent threats in the modern landscape. Invest in tools, technologies, and training designed to detect and defend against the full spectrum of deepfake-related incidents.

Section 4

Training Is Widespread but Its Efficacy Is Questionable

The overall adoption rate and frequency of deepfake-focused security awareness training is increasing. In the present study, over **88% of respondents said their organizations have provided deepfake-related cybersecurity training, up from just 68% in 2024.** Meanwhile, just over 1 in 10 (11%) say they have never participated in deepfake-related cybersecurity training.

Of equal importance is the frequency with which organizations are providing those trainings. Of those respondents that had received training, the largest proportion (at 44%) did so quarterly. Encouragingly, well over a third (38%) of respondents said they were receiving monthly deepfake-related training—all increases from prior years' results.

Q6. If so, how often does your organization provide such training?

*Those whose organization has provided cybersecurity training related to identifying deepfakes

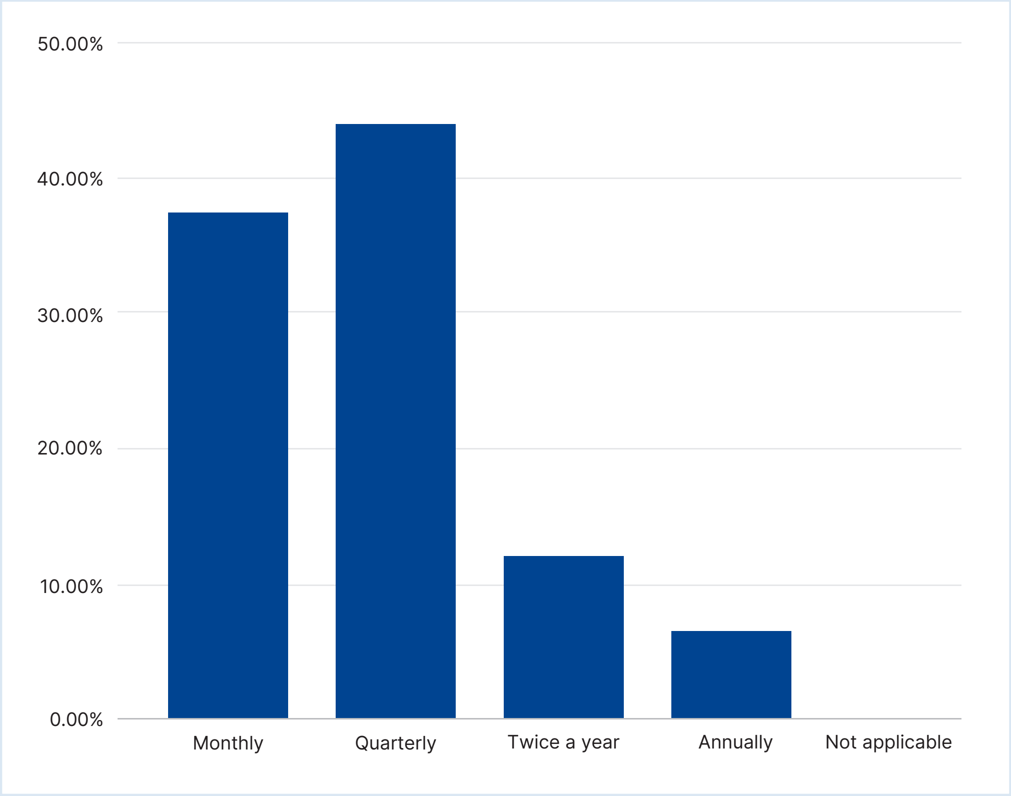

88%
of respondents said their organizations have provided deepfake-related cybersecurity training, up from just 68% in 2024.

The surrounding findings from this year's study suggest that the efficacy of these training remains limited. When looking at organizations' performance in deepfake-related security testing, the results were less than stellar:

- Only 8.4% of organizations saw first-try simulation pass rates of 80% or above.
- Nearly three-quarters (74%) of organizations saw simulation pass rates below 60%.
- The first-try deepfake-simulation pass rate average was 44%, suggesting a significant skills/retention gap worth addressing.

The real-world data reinforces questions about training efficacy.
Despite widespread training adoption, 94% of respondents remain concerned about deepfake threats, 85% continue to experience incidents, and 55% report financial losses in the past year, frequently in substantial amounts.

In both simulations and real-world outcomes, employees remain susceptible to deepfake-related incidents, meaning there is still much work to be done to improve employee awareness and vigilance (especially as deepfaked media continues to grow more realistic/sophisticated). Organizations must ensure that their training and testing tools employ AI and real-world analysis so that training and simulations keep up with the changing threat landscape.

Despite widespread training adoption,

94%

of respondents remain concerned about deepfake threats.

Perhaps even more importantly, organizations should be ready to double-down on tools and technologies designed specifically for deepfake defense. The human element still plays a primary role in roughly 60% of all cyberattacks, and deepfake-related incidents are no exception. With this in mind, it's critical that organizations continue investing in training, and couple that with next-generation AI-powered defensive technologies, and deepfake-detection tools to lessen the impact of these increasingly common and costly attacks.

As discussed in Section 1, the rising scale of financial losses is more than enough to assume a more than adequate ROI when investing in new, AI-powered cybersecurity technologies. While training remains absolutely essential for organizational security, it's increasingly important that those organizations supplement training with technologies and policies that limit the scope and damage done by human error alone.

Section 5

While Concern Skyrockets, Defensive Overconfidence Grows

When asking IT/security professionals whether or not they were concerned about deepfake-related incidents, the response was a resounding “yes.” In this year’s study, nearly everyone (94%) expressed at least some level of concern about the threat deepfakes currently pose to their organizations. Meanwhile, over 63% reported being “very concerned” about the threat they pose, a more than 15% increase from last year’s findings—indicating an increasing severity of concern around deepfake-related incidents.

Interestingly, these percentages changed very little (with average increases of just a few percentage points) when the question was reframed to ask how concerned they are about the threat in the near future (with net concern, “very concerned,” and “somewhat concerned” sitting at 96%, 67%, and 31%, respectively). This seems to reinforce the idea that deepfake-related incidents are no longer being viewed as some emerging threat looming just over the horizon. On the contrary, it is a real and present danger to security teams and their organizations; and IT professionals know it.

94%
expressed at least some level of concern about the threat deepfakes currently pose to their organizaions.

Q1. When considering deepfake threats in general (image, video, voice, email), how would you rate your concern about the risks to your organization both now and in the future?

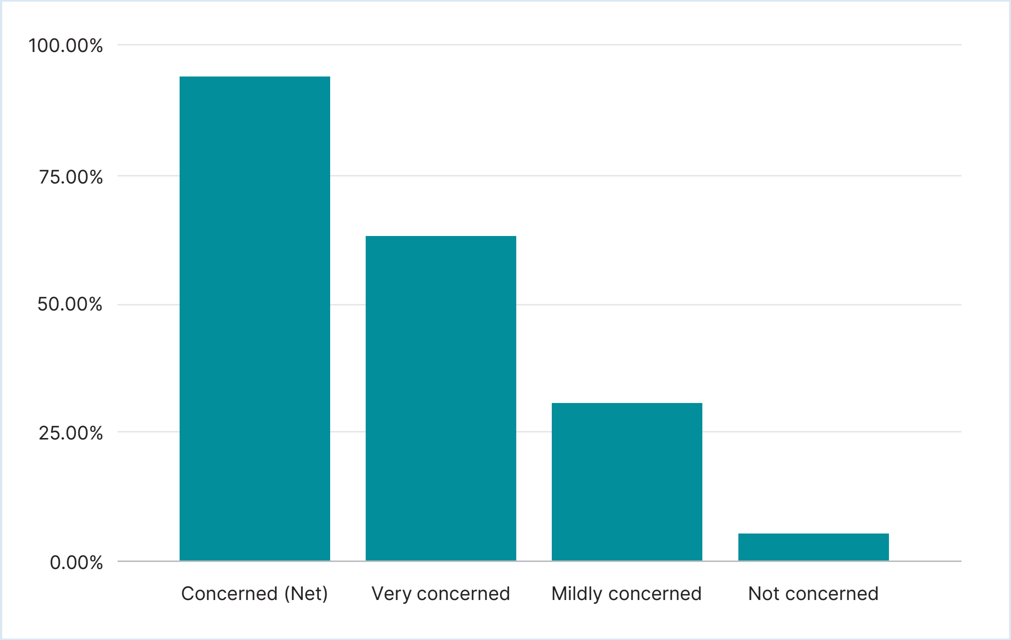

The survey showed 99% of respondents claim to be confident in their organizations' ability to defend against deepfake-related incidents. And the degree of that confidence is growing. The percentage of respondents saying they were "very confident" has shot up by over 30% in just one year, from just over 40% in 2024 all the way to 73% in 2025. Meanwhile, the percentage saying they were "mildly or somewhat concerned" fell from over 52% in 2024, to just 31% in 2025.

This disconnect might be dismissable as just a matter of question design if organizations were otherwise performing well in defending against these threats. On the contrary, though, as we've seen throughout this report, organizations are not faring well thus far in the fight against deepfakes. The overwhelming majority of respondents have experienced one or more incidents in the past year, and the majority of them have suffered significant financial losses as a result (at a mean cost of over $280,000).

So, what's behind this seeming preparedness paradox? Some possible explanations include professionals' wishful thinking about their organizations investing more and improving their defensive capabilities, an underestimation of more widespread prevalence, or plain old wishful thinking.

99%
of respondents claim to be confident in their organizations' ability to defend against deepfake-related incidents.

Conclusion: Guidance for Organizations Facing the Deluge of Deepfake-Driven Threats

Deepfake-related incidents are no longer futuristic warnings, they are a present and escalating threat with material financial consequences. The growing volume, sophistication, and multi-modal spread of these attacks demand more than acknowledgment; they require concrete prioritization and investment.

To effectively combat these threats, organizations should focus on three critical areas: training that addresses the full spectrum of deepfake sophistication, detection technologies that match the pace of AI advancement, and incident response processes that account for the unique challenges of synthetic media.

Ultimately, preparedness against deepfake attacks requires continuous evolution of defenses. Early investments in comprehensive strategies will determine which organizations can effectively defend against this rapidly evolving threat.

Unsure where to begin? Start by reaching out to the experts at IRONSCALES and discover how we can help you bolster your deepfake defenses.

Methodology

This original research was commissioned by IRONSCALES and conducted by Censuswide. The survey polled a sample of 500 US-based IT/Cybersecurity professionals (manager level to the C-Suite). All participants were ages 25-55 (all genders) and worked at companies of 1,000-10,000 employees. The data was collected between 09/08/2025 and 09/16/2025. Censuswide abides by and employs members of the Market Research Society and follows the MRS code of conduct and ESOMAR principles. Censuswide is also a member of the British Polling Council.

About IRONSCALES

IRONSCALES is the leader in AI-powered email security protecting over 17,000 global organizations from advanced phishing threats. As the pioneer of adaptive AI, we detect and remediate attacks like business email compromise (BEC), account takeovers (ATO), and zero-days that other solutions miss. By combining the power of AI and continuous human insights, we safeguard inboxes, unburden IT teams, and turn employees into a vital part of cyber defense across enterprises and managed service providers. IRONSCALES is headquartered in Atlanta, Georgia. To learn more, visit www.ironscales.com or follow us on LinkedIn.