

The Deepfake Readiness Index:

Healthcare edition

How security leaders think about
their own deepfake defenses





73%

of healthcare organizations report or suspect deepfake attacks, but **only 9%** have tools built to catch them.

Deepfakes are more than a nuisance or an isolated issue for hiring or PR teams. Deepfakes are squarely a security issue, one that can lead directly to data theft and intrusion.

Most companies have already adopted tools that model Zero Trust, embedding verification across their network and systems. Yet deepfakes evade those controls; they can pose a risk to every virtual business interaction: phone calls into your helpdesk, remote job interviews, and sensitive video meetings with executives. Until companies treat deepfakes like the serious enterprise threat they are, most will remain vulnerable.

Pindrop surveyed 122 U.S. healthcare security leaders to find out exactly how unprepared companies are.

94%

are concerned they aren't prepared for a deepfake attack.

72%

believe deepfake defenses will improve faster than attack quality.

~1 in 4

affected organizations lost \$1 million or more to a single deepfake incident.

70%

of leaders say it will take an impersonation of a company leader to make this a boardroom priority.

30%


believe one successful attack could put a company like theirs out of business.



38% of healthcare orgs are fighting deepfakes with obsolete, legacy defenses

Security leaders report that they're concerned about deepfake attacks (94%), yet their organizations aren't rolling out tools that directly address this concern.

What tools are you using to combat deepfake identity attacks?



Only 9% of respondents had purpose-built tools for the deepfake threat.

Legacy security tools (IAM, SSO, MFA) are no match for the scale and sophistication of today's deepfake attacks, largely because attacks are happening in channels that these tools weren't designed to defend. Additionally, humans working in those channels are simply not able to detect deepfake and synthetic content reliably.

What percentage of your company's workforce do you believe could identify a deepfake attack, such as a fake customer, vendor, or someone impersonating a company leader?



87% of respondents believe half or less of their workforce could identify a deepfake attack. When combined with separate research that suggests humans are not reliable in detecting synthetic content,² the hope of depending on human judgment to stop deepfake attacks is risky.

Real-time communication channels have become a prime target for exactly these reasons. Without a reliable way to evaluate whether the person on a video meeting is real, organizations are flying into a danger zone with zero visibility and attackers know it.

In February 2026, Pindrop reported on the spike in AI attacks. At that time, among the Pindrop customer base, from Q4 2024 to Q4 2025, AI attacks had jumped by **1210%**. Just three months later, that number was **1390%**. And by the end of June 2026, the growth had reached **1680%**. Compared to traditional attack types, AI attacks are **growing 8x faster**.³

Deepfake detection point-solutions may seem like a promising answer to the crisis. But a point solution is not a long-term solve. AI has created a **foundational identity and trust question** that requires an equally foundational change in the security strategy for real-time channels. Those channels have now become identity checkpoints where eyes and ears are no longer trustworthy. This is why a new layer of defense is needed, one that extends *continuous, real-time* Zero Trust to the channels organizations didn't used to worry about.

“

For thousands of years, trust was based on perception: recognizing someone's voice or seeing someone face-to-face. We're now in an age where trust can be synthesized. AI has changed what we think of as proof.”


Ajit Gaddam

Head of Fraud, Financial Crimes, and Trust Systems,
HealthEquity

The survey reveals a notable disconnect. While 72% of leaders are optimistic that security defenses will outpace the quality of attacks, only 9% have invested in purpose-built tools.

Another survey question reveals that this is a *prioritization problem*, with 46% of respondents saying their leadership has not prioritized investment or does not plan to increase investment in defending against AI-driven attacks.

Which of these best describes how your organization's leadership has prioritized defending against AI-driven attacks?



PINDROP POV

While there is widespread concern and confidence in future technological progress, that sentiment has yet to translate into widespread immediate investment in specific deepfake defenses.

There's real exposure, but this attack surface appears to not be prioritized in the boardroom yet.


But for the 54% who have invested and plan to increase investment, there is a clear signal that healthcare leaders are seeing the risk and plan to counter it. Of those that have prioritized and invested, **52% report a strong return on investment**.



A single deepfake is a million-dollar security failure


When a deepfake attack lands, the financial exposure is significant. Among organizations that experienced or suspected a deepfake attack, **42% reported total costs of \$500,000 or more**. 23% crossed the \$1 million threshold.

Total cost of attack



Includes direct losses, remediation, and staff time.

Consequences that followed



% of affected organizations reporting each consequence

What came next only added to the damage. The top three consequences were:

1. 47% of affected organizations faced a data breach
2. 44% faced a compliance risk
3. 40% experienced a follow-on cyberattack like ransomware

These consequences carry immense weight, particularly when considering the regulatory and operational fallout of a PHI breach. When a breach occurs, whether through a deepfake-enabled social engineering attack or otherwise, organizations are exposed to severe regulatory penalties and potentially significant operational paralysis. The infamous [2024 Change Healthcare breach](#) is a definitive example in how credential compromise, often initiated via social engineering, can cascade into system-wide failure.

More recent examples ([DentaQuest, 2026](#); [Boston Scientific, 2026](#); [Xsolis, 2026](#); [NYC Health + Hospitals, 2026](#)) on the [ever-growing list of healthcare breaches](#) continue to heighten the pressure to tighten security across all entry points in order to protect patients, data, and reputation.

In May 2026, Pindrop [published a report](#) based on findings from the CISO Deepfake Defense Council, a select group of security executives from seven Fortune 500 enterprises and other category-leading organizations. The Council is focused on delivering strategic guidance to help enterprises understand, prepare for, and defend against deepfakes, one of the fastest-emerging threats to trust and identity.


Council members identified and mapped five AI attacks, based on reputational and financial risk. Three of those attacks (fake job candidates, IT helpdesk attacks, and executive impersonation) landed as catastrophic or critical threats, reflecting high risk both financially and reputationally.⁴

That's why the missing detection layer matters. The financial and reputational exposure from a deepfake attack compounds with every hour it goes undetected. If security leaders understand the problem and know that the implications can be severe, why is the adoption of purpose-built tools so low?



Deepfakes won't become a boardroom priority until one hits an executive

How strongly do you agree or disagree with the following statement: "It will take a leader at my organization (executive, HR, or finance leader, etc.) being personally fooled or impersonated by a deepfake for them to take this threat as seriously as they should."



70% of security leaders agree it will take a leader at their organization being *personally impersonated* before deepfakes become a genuine boardroom priority. That's **7 out of 10 security professionals**, people who understand the threat most clearly, who don't believe their leadership teams will act until the threat hits home.

It's possible boardrooms still view deepfakes as a public figure or celebrity issue, one defined by manipulated videos on social media that spread like wildfire and make headlines. That framing made it easy for leadership teams to file deepfakes under "reputational risk for high-profile individuals" rather than "persistent enterprise security threat requiring infrastructure investment." It's also true that healthcare organizations face much more than reputational damage. Attackers repeatedly target patient data, day-to-day operations, and the financial integrity of all types of healthcare organizations.

PINDROP POV

Deepfakes are today's identity attacks, and need to be treated as such. This means recognizing that the live communication channel is a gap in an otherwise Zero Trust stack, and closing that gap requires the same commitment organizations brought to endpoint protection, cloud security, and identity management.

The next step is getting the board and higher-level leadership to understand the risk in terms they already use: a threat that is active, financially significant, and getting more sophisticated faster than the defenses already in place.



The security leaders prioritizing deepfake detection today are the only ones who'll have a real answer when their board starts to ask questions."


Jim Routh

Chairman of The CISO Deepfake Defense Council and former 6X CSO/CISO in financial services and healthcare



Nearly a third of healthcare leaders think deepfakes are a threat to a company's survival

"How likely do you think it is that the damage caused by a single deepfake attack could cause a company like yours to go out of business?"



30% of security leaders consider it at least somewhat likely that a single deepfake attack could put a company like theirs out of business, including 16% who consider it extremely likely or certain. Meanwhile, 70% think it's unlikely a deepfake attack would cause a company to go out of business.

Even though 70% don't fear their company going out of business due to a deepfake attack, deepfake attacks are still costly, in the hundreds of thousands, even millions, as seen in Finding 2. For the 30%, who see a successful synthetic identity attack as an existential threat, the financial and indirect costs could be severe enough to close their doors.

PINDROP POV

Deepfakes are an identity attack that exploits how your organization verifies trust in human communication. This attack can touch **every function in your business**: patients or members calling to access personal medical or financial information, IT resetting credentials, HR interviewing and hiring new staff, an executive making a decision based on a voice they recognized. The impact of a single successful attack isn't contained to one system or one team.




FINDING 4 [CONTINUED]

Greater than 40% of healthcare security leader respondents say their company has experienced IT helpdesk deepfakes and executive impersonations; 38% of respondents believe they have experienced call center attacks; and 28% have experienced deepfake job candidates.

To get the board to prioritize protections correctly, it's important to educate them on how these attacks permeate the enterprise and why they are more than just high-profile news headlines.

Which of the following deepfakes do you know or suspect your company encountered in the past 12 months?





CONCLUSION

Pindrop: your deepfake detection partner

AI has introduced doubt into every real-time communication, which has left many security leaders at a crossroads: do you patch your current tech stack or do you rethink your identity strategy entirely?

Since 2010, Pindrop technology has defended Fortune 500 companies and their patients, members, and customers against voice fraud, identity attacks, and social engineering, including the top 3 healthcare companies, 14 of the top 20 banks, and 5 of the top 7 insurers. Across more than 8 billion analyzed interactions, we became experts in fraud and attack tactics across voice and video channels.

Pindrop technology is built on three trust engines: synthetic detection, risk analysis, and authentication. The Pindrop platform brings together all three in real time, verifying voice and video in virtual meetings and other real-time interactions.

So your organization has the trust it needs, the moment it needs it.

Where does your organization land?

[TAKE THE ASSESSMENT](#)[TALK TO A REAL HUMAN](#)

Sources

¹ The Pindrop Survey was conducted by Wakefield Research (www.wakefieldresearch.com) among 250 US Security Leaders at organizations with a minimum of 1,000 employees, between June 11th and June 22nd, 2026, using an email invitation and an online survey. Results of any sample are subject to sampling variation. The magnitude of the variation is measurable and is affected by the number of interviews and the level of the percentages expressing the results. For the interviews conducted in this particular study, the chances are 95 in 100 that a survey result does not vary, plus or minus, by more than 6.2 percentage points from the result that would be obtained if interviews had been conducted with all persons in the universe represented by the sample.

Unless otherwise noted, all figures in this report are drawn from the 122 respondents at healthcare organizations within that sample. Subgroup findings carry a wider margin of error than the total sample: ± 8.9 percentage points at 95% confidence for the healthcare subset ($n=122$). Wakefield notes that base sizes under 100 are directional findings only.

² Several studies describe the inability of humans to detect synthetic content reliably. Cooke et al., "As Good As A Coin Toss," 2025, found humans detect AI-generated content correctly only about 50% of the time, no better than a coin flip. Barrington et al., "People are poorly equipped to detect AI-powered voice clones," 2025, concluded that human listeners could correctly identify a voice as AI-generated or real only about 60% of the time. A third study, Müller and Choong, "Eroding Trust in Real Speech: A Large-Scale Study of Human Audio Deepfake Perception," 2026, found people have grown more skeptical of all audio, and accuracy on genuine recordings actually dropped by 8.6 percentage points relative to a 2021 baseline.

³ Based on Pindrop customer data analysis from Q2 2026 in comparison to the previous six quarters.

⁴ Compiled based on five anonymous mapping exercises with security leaders from the CISO Deepfake Defense Council.