Pindrop

The 5 high-cost AI attacks every security leader needs to know.

Your trusted channels have a blind spot: **AI-backed attacks**.
Find out where you're exposed.

Here's what AI attacks actually look like

Discussions about AI-powered attacks happen behind closed doors, tucked away in internal incident reports and confined to candid exchanges between peers at industry conferences. This secrecy leaves security leaders with an near-impossible task: **responding to threats they can't see.**

That's why we're breaking down the anatomy of AI attacks: what they look like, how they work, and why they're able to slip past traditional defenses.

Here are the stakes: A multi-million dollar wire transfer after a deepfake video call.[1] A North Korean operative moving through your hiring pipeline.[2] An IT helpdesk handing over SSO credentials to someone posing as an employee.[3]

Attacks like this are playing out at companies everywhere, everyday. And with **AI-driven attacks growing by nearly 14x in six quarters**,[4] organizations can't afford to treat this as low-risk threat.

Citations

[1] CNN, "Finance worker pays out $25 million after video call with deepfake 'chief financial officer'," February 2024.
[2] Pindrop, "From Interview to Intel Drop: The Moment We Exposed a Coordinated Hiring Scheme," July 2025
[3] Pindrop, "How Pindrop Technology Could've Prevented the MGM Breach," September 2023.
[4] Based on Pindrop customer data from Q4 2024-Q1 2026. Derived from a study of over 700M calls.

AI has forced security leaders to doubt the channels they trust

Talk to a CISO right now and one word keeps coming up: uncertainty. Uncertainty about how AI attacks will reshape security strategies. Uncertainty about where to invest. Uncertainty about where the real risk even lives.

That's exactly why we created the **CISO Deepfake Defense Council.** The Council brings together a select group of highly regarded security executives from seven Fortune 500 enterprises and other category-leading organizations across a broad cross-section of industries, including technology, financial services, healthcare, life sciences, defense, enterprise software, cryptocurrency, and cybersecurity. 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.

We interviewed these leaders individually and as a group, probing to understand how they're thinking about the new enterprise threat landscape.

Takeaways from the CISO Deepfake Defense Council

Unauthorized access is just one step in a longer attack chain.

Impersonation (deepfakes or otherwise) opens the door. Once a bad actor gets in, they become an insider threat, often with the access they need to cause system-wide damage.

Telling the difference between a good bot and bad bot is about to be critical.

As companies incorporate AI into their workflows and authorized agents begin handling tasks, companies will need a way to discern what is an approved use of AI and what is suspicious, unauthorized activity.

Attackers are playing in blindspots.

It's nearly impossible to defend against what you can't see. That's why bad actors attack the channels that seem low-risk or protected (e.g., hiring or vendor comms). Without visibility, leaders can struggle to calculate exposure.

Perceived weak points vary across industries.

One security leader points to helpdesk impersonation as a critical risk, another points to contact center attacks. This variation seems to be driven by industry-specific priorities. For example, financial services may prioritize high-risk customer or client interactions that can end in direct financial losses. Leaders from other industries may perceive different interactions as more vulnerable.

Companies need detection that works with them, not against them.

Adoption of deepfake detection needs to integrate seamlessly into existing workflows. Roadblocks will directly impact whether people utilize the tool when its needed.

How CISOs think about the new threat landscape

Council members mapped AI attacks by reputational and financial risk. Answers varied, but six threats rose to the top.[1] This briefing covers those six.

> A healthy immune system needs to be proactive, anticipating threats before they attack. In the age of deepfakes, enterprises need to approach defense in the same way: building their cybersecurity strategy for resilience and adaptability.

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

Citations

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

Where do these attacks land on the risk exposure grid for your organization? Download a blank version and bring this exercise back to your team.

FAKE JOB CANDIDATES

How AI helps fake job candidates get hired

Attackers use synthetic identities and real-time deepfakes that help them pass interviews to land on your payroll.

Your applicant pool is full of potential insider threats

In 2025, the FBI released an official public service announcement warning about North Korean operatives trying to get hired at U.S.-based companies.[1] These operatives use fabricated or stolen identities to apply for jobs, funneling salaries back to the regime while conducting espionage from inside company walls.

Deepfake applicants show up with polished resumes, credible LinkedIn® profiles,[2] and impressive work history, and when they log on for an interview, they seem well-prepared and professional.

But these candidates often use real-time face-swap tools trained on photos scraped from real people's social media, filtered or manipulated voices, and work references that lead to companies that don't exist or can't be verified.

And the consequences of hiring a fake candidate can be **severe.** Once inside a company's systems, a fraudulent employee can steal data and intellectual property, or plant backdoors for future exploitation.

Companies that unknowingly pay these workers may also find themselves in violation of federal sanctions laws. And being infiltrated by a state-sponsored actor can erode customer trust and bring unwanted scrutiny.

1 in 343

job applicants were linked to North Korea in Pindrop's study **in 2025.**[3]

the next year...

1 in 47

job applicants were linked to North Korea in Pindrop's study **in 2026.**[4]

resulting in...

630% increase

in DPRK-linked job applicants for Pindrop in ***just one year.***[5]

Citations

[1] FBI, "North Korean IT Worker Threats to U.S. Businesses," July 2025.
[2] LinkedIn is a registered trademark of LinkedIn Corporation and its affiliates in the United States and/or other countries.
[3] Pindrop, "From Interview to Intel Drop: The Moment We Exposed a Coordinated Hiring Scheme," July 2025. Based on Pindrop's 2025 internal investigation of applicants for two remote roles.
[4] Based on Pindrop's 2026 internal investigation of applicants for a Senior Software Engineer role.
[5] Comparison of Pindrop's 2025 internal investigation and 2026 internal investigation.

REAL WORLD CASE STUDY

“Jamie” was likely a North Korean operative

In early 2025, Pindrop started investigating fake job candidates. Since then, the deepfakes have only gotten more convincing. Case in point: “Jamie,” a candidate for Pindrop’s Senior Software Engineer role.

Jamie had clear answers, a polished delivery, and seemed well-prepared. It took a Pindrop Pulse® for meetings alert to raise concerns. Without it, he likely would have moved forward. Pindrop later linked Jamie’s IP address to North Korea.

Pindrop researchers also noticed patterns in device telemetry, geography and network characteristics, email patterns, and synthetic identity construction amongst fake candidates. By extracting attributes from a confirmed fraudulent candidate, they built a relational graph of past applicants—and found that **one confirmed fake applicant helped uncover a web of 23 past candidates with similar patterns.**

One fake candidate. A whole web of fraudulent applicants.

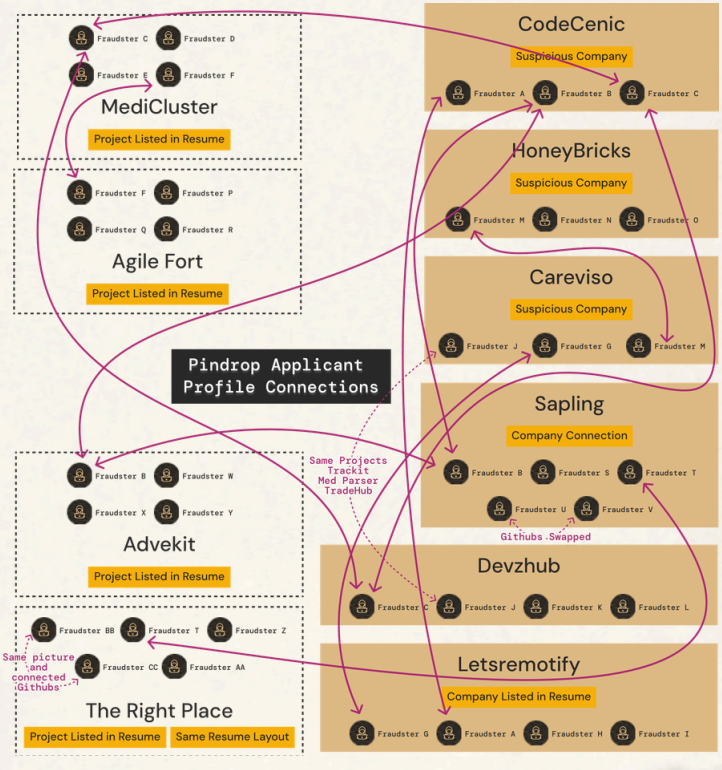

Don’t wait for a fake job candidate to slip past your defenses

Take action today. **Download the Deepfake Readiness Checklist.**

CONTACT CENTER ATTACKS

Attackers continue to exploit trust in contact centers

AI helps attackers automate their schemes, speeding up account takeovers and compounding losses.

How the manipulation scam works

This sharp rise in AI-driven attacks is happening in contact centers, where companies often still rely on knowledge-based authentication questions to verify identity.

The attack chain is simple: impersonate a patient or member, pass the authentication check, steal PHI, or even take over an account. And authentication is often easy to pass with data from breaches. Now with AI, these attackers can increase their operations exponentially, inundating contact centers with a high volume of attacks.

And when it works, the damage hits fast. PHI and sensitive data is released, payments are rerouted, funds are drained from accounts. For the organization, that means compliance risk, fraud losses, costly incident reviews, and reputational damage. In an industry where trust is everything, a wave of contact center attacks can have lasting consequences.

1390% increase

in AI attacks in just six quarters.[1]

60x growth

in AI attacks per day[1]

An average 287 AI attacks per day across our customer base, a staggering **6400% increase from 2025**.

8.7% of all fraud

is AI-backed in contact centers as of Q1 2026[1]

This is up from the prior year, when <1% of fraud in the contact center was non-live.

Bot calls are getting more realistic fast.
Listen to these real-world calls from the past year. Can you tell the difference?

Spring 2025 | Summer 2025 | Fall 2025

Citations

[1] Based on Pindrop customer data analysis from Q1 2026 in comparison to the previous five quarters.

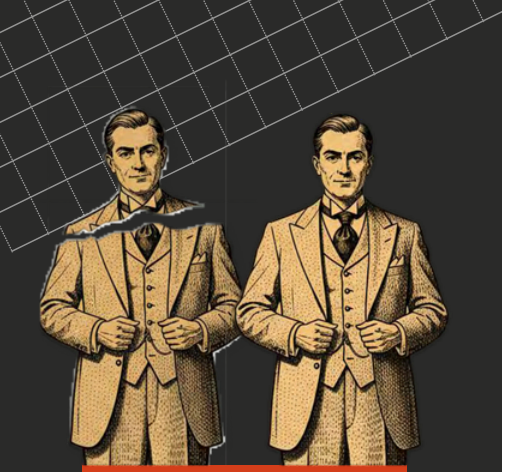

EXECUTIVE IMPERSONATION

Your exec team is an easy impersonation target

Attackers use AI to study and replicate executives' voices and faces—creating convincing clones in seconds.

Weaponizing authority for deception

When an employee gets a suspicious email, security training tells them to verify through a different channel. But when they join a video call and see a face they recognize, that feels like verification. Even though it isn't. In the AI era, seeing and hearing someone is no longer proof that they're real, and most employees don't know that yet.

Executive impersonation is happening now, and virtual meetings are one of the biggest blind spots enterprises have. A cloned voice or face in a meeting carries the same authority as the real thing, and in many cases, nobody in that room can tell the difference. A random text asking for gift cards is easy to dismiss, but a call that sounds and looks identical to your CEO? **That's much harder to ignore.**

REAL WORLD EXAMPLE

According to a 2025 PSA from the FBI,[1] bad actors have been impersonating government officials, including Cabinet members and members of Congress, to gain unauthorized access since 2023. These attackers send AI-generated voice messages and texts, posing as senior U.S. officials to build rapport with targets before manipulating them for their own gain.

In the summer of 2025, an AI-generated voice clone impersonating Secretary of State Marco Rubio successfully reached senior U.S. officials and foreign ministers.[2] Just a synthetic voice lifelike enough to pass as a high-ranking government official.

These same tactics are already being used against corporations. Anyone in a position of authority or influence is an impersonation target. Fake CEOs releasing financial documents, making public announcements, manipulating an employee—all can lead to serious reputational, operational, and regulatory fallout.

Citations

[1] FBI, "Senior U.S. Officials Continue To Be Impersonated in Malicious Messaging Campaign," July 2025.

[2] Pindrop, "AI Voice Deepfakes Fooled World Leaders. What Happens When They Target Your Employees Next?," July 2025.

Deepfake leaders can end in multi-million dollar mistakes

Think about what happens when an employee joins a meeting with their CEO who asks them to move fast on a wire transfer. The instinct is **to help**, to **not slow things down**, and to **trust the person at the top**. Executive impersonation works because it exploits something security training struggles to override: the instinct to comply with authority.

And the damage goes beyond manipulated employees. Fake videos of leaders making statements about financials or product updates can do serious damage to brand reputation too. Leaders have access, influence, and trust. AI lets scammers mimic all three.

We created a deepfake of Pindrop's CFO to prove just how easily attackers can join meetings as a fake executive.

REAL WORLD EXAMPLE

Attackers impersonated WPP's CEO using a fake WhatsApp account, cloned voice, and deepfake video to pressure an executive for money and sensitive information.[3]

REAL WORLD EXAMPLE

Scammers ran a highly realistic Elon Musk deepfake on YouTube to promote a crypto scam. It pulled in 140,000 viewers and stayed live for hours.[4]

REAL WORLD EXAMPLE

In May 2026, The Singapore Police Force released a statement about a scam involving a fake Zoom meeting impersonating Singapore's Prime Minister and other senior government officials to solicit phony government funding. One victim wired S$4.9 million ($3.8 million USD) before the scheme was uncovered.[6]

REAL WORLD EXAMPLE

Arup lost $25 million after an employee joined a video call with what looked and sounded like their CFO. It was a deepfake, and the transfer went through before anyone caught it.[5]

Citations

[3] The Guardian, "CEO of world's biggest ad firm targeted by deepfake scam," May 2024.
[4] Pindrop, "Elon Musk Deepfake Exposes AI-Generated Fraud Dangers," July 2024.
[5] CNN, "Finance worker pays out $25 million after video call with deepfake 'chief financial officer'," February 2024.
[6] South China Morning Post, "How a victim lost US$3.8 million in Singapore deepfake Zoom scam impersonating PM Wong," May 2026.

IT HELPDESK ATTACK

IT helpdesks are easy to overlook as a target

And that's exactly why bad actors attack there. With one MFA reset, they can take down your entire system.

How helpdesk attacks play out

Armed with basic personal data from social media and leaked credentials, most of it available for mere dollars on the dark web, attackers call in, impersonate an employee, and request an MFA reset. If the helpdesk relies on knowledge-based authentication alone, they open the door for the scammer. From there, it's a straight path to unauthorized access, ransom, IP theft, and the ability to move anywhere in your systems.

That's how the Scattered Spider attack against MGM Resorts played out.[1] The hackers found an MGM employee's LinkedIn® profile,[2] impersonated that employee on a call with the helpdesk, and asked for login support. From there, the hackers gained access to MGM's internal systems.[1]

The damage was severe. Widespread system outage impacted systems for several days.[3] Caesars Entertainment and MGM shares fell in the immediate aftermath3 and the organization faced a class action lawsuit over the protection of customer privacy.[4]

REAL WORLD EXAMPLE

Microsoft released a warning to employees in 2026:[5] scammers are infiltrating organizations, sending Teams messages pretending to be IT or helpdesk staff, and requesting remote access to victims' computers. Once inside, they get access to sensitive data, systems, and IP.

"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."

Head of Fraud, Financial Crimes, and Trust Systems at HealthEquity
Ajit Gaddam

Citations

[1] University of Hawaii—West O'ahu, "ALPHV: Hackers Reveal Details of MGM Cyber Attack," October 2023.
[2] LinkedIn is a registered trademark of LinkedIn Corporation and its affiliates in the United States and/or other countries.
[3] Reuters, "MGM Resorts breached by 'Scattered Spider' hackers: sources," September 2023.
[4] Fox5 Vegas, "MGM begins payouts in $45M data breach settlement," December 2025.
[5] Tech Radar, "Microsoft issues warning over Teams helpdesk impersonation attacks – hackers are 'blending into routine IT support activity' by abusing remote assistance access," April 2026.

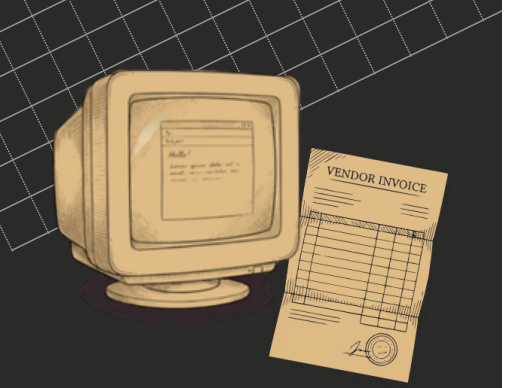

VENDOR OR PARTNER IMPERSONATION

How a vendor relationship becomes a vulnerability

External partners, vendors, and suppliers are often vital for a company's operations, which is exactly why attackers impersonate them.

Third-party relationships are invaluable, but easily manipulated

Vendor impersonation and partner impersonation are some of the **most underestimated risks** enterprises face right now. These are relationships built on trust, and they operate across external systems that are inherently harder to monitor and verify. Unlike internal communications, the channels aren't as secure. An attacker posing as a trusted supplier just needs to send a realistic email, a cloned voice message, or a fake invoice to break through defenses.

A financial institution (and Pindrop customer) believed initially they only had a problem with deepfakes in hiring. After only deploying Pindrop Pulse® for meetings to their hiring team, an employee outside their recruiting team flagged a supposed "vendor" as suspicious. They asked for access to Pulse and within seconds, caught a deepfake.[1]

REAL WORLD EXAMPLE

A cybersecurity company reported attackers using AI to create realistic email threads and fake invoices that looked legitimate. Payments were approved and sent before anyone realized the vendor wasn't real.[2]

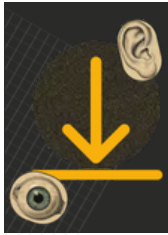

Humans reliably catch AI manipulation only ~50% of the time[3]

A coin toss. That's how good humans are at catching deepfakes. Technology purpose-built for detection is just better at spotting deepfakes than humans. This tech looks at the subtle qualities like resonance, timbre, and facial movements that separate human media from AI. Read the datasheet to learn how it works.

Download the datasheet

Citations

[1] Anonymous Pindrop Pulse for meetings customer.

[2] Mimecast, "BEC Campaign Using AI Generated Fake Email Threads," August 2025.

[3] Cooke, Di, et al. "As Good as a Coin Toss: Human Detection of AI-Generated Content." Communications of the ACM, edited by , vol. 68, no. 10, Sept. 2025, pp. 100–9. Crossref, https://doi.org/10.1145/3729417.

The question used to be:
is this a fraudster?

Now its:
is this a real human?

For attackers, AI is becoming the force behind their operations

A single check at login means nothing when attackers can generate a convincing identity, voice, and face in real time, and enter the channels you trust. The answer is **continuous identity verification:** ongoing, multi-signal validation throughout every real-time interaction, every call, every remote interview, every vendor communication.

The leaders who've been hit already know this.

The ones who haven't are running out of time to get ahead of it. The cost of waiting is a multi-million dollar wire transfer, a fraudulent new hire, or a breach that starts at your IT helpdesk.

Be proactive. Talk to an AI identity expert. →

REAL WORLD PRODUCT HERE TO HELP

Pindrop® tech helps defend against PHI breaches, helpdesk attacks, bot calls, and more

Recognized by **TIME as one of the Top 10 Most Influential Software Companies of 2026**, Pindrop solutions deliver synthetic voice and deepfake detection, participant authentication, and location intelligence across leading CCaaS and conferencing platforms, defending critical patient and business conversations from AI and human impersonation with industry-leading accuracy.

How the Pindrop® platform is different

- **Real human**: Highly accurate detection of deepfakes powered by industry-leading liveness models across audio and video.
- **Right human**: Only platform that passively authenticates to confirm not only that a speaker is human, but that they re the right person.
- **Right place**: Only platform that provides location intelligence, flagging risk indicators such as VPN usage or mismatched geographies.

Seamlessly integrated into major CCaaS platforms, like NiCE, Genesys, and Five9,* and meeting platforms, like Zoom,[1] Webex,[2] and Microsoft Teams,[3] Pindrop solutions surface real-time risk signals to help you interact confidently with the right callers and meeting attendees.

Pindrop

Citations

[1] Zoom and the Zoom logo are trademarks of Zoom Communications, Inc., registered in the U.S. and other countries.

[2] Cisco, Cisco Webex, and Webex are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

[3] Microsoft Teams is a trademark of the Microsoft group of companies.

*visit https://www.pindrop.com/partner/ to learn more about our partners