Adaptive

The Rise of SMS and Voice Phishing

How Smishing and Vishing Became Today's Most Dangerous Attack

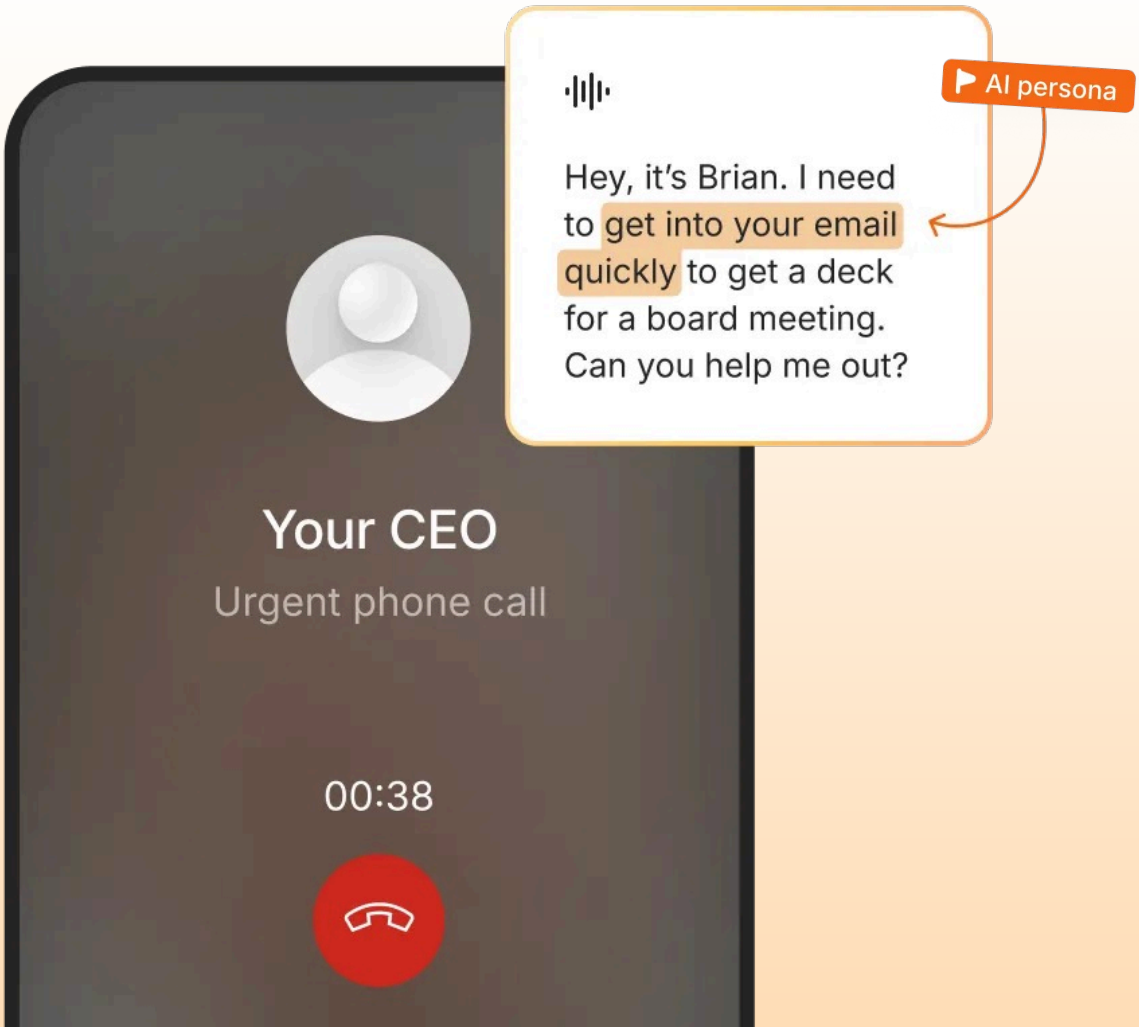

Table of contents

The Breach Didn’t Begin With Malware. It Began With A Phone Call.

A threat actor contacted the company’s IT helpdesk, claimed they were an employee who had lost access, and persuaded support staff to reset credentials. That single, human interaction opened the door to a wider compromise—one that eventually disrupted operations, exposed sensitive systems, and triggered a costly incident response across one of the largest hospitality brands in the world.

People Become The Primary Target.

That breach was MGM Resorts, in September 2023. The attack was widely attributed to a social-engineering-focused group known as “Scattered Spider,” which specializes in impersonation and helpdesk manipulation rather than traditional technical exploitation. And the key lesson wasn’t technical sophistication—it was channel selection. The attacker didn’t force their way through hardened email defenses. They moved around them, exploiting the pathways security teams monitor the least: phone calls, SMS messages, and mobile-first workflows where trust is assumed and verification often breaks down.

This is what modern social engineering looks like: it doesn’t break systems—it gets people to break process. Smishing and vishing have moved from fringe tactics to primary attack paths because attackers have learned where defenses are weakest and trust is strongest.

This paper explains how that shift unfolded: how SMS and voice channels bypass technical controls, how generative AI has changed the scale and believability of impersonation, and why manipulation now happens far beyond the inbox. It also outlines what organizations need to do next—because the decisive moment is no longer a click, but a conversation.

As attacks move into real-time, mobile-first interactions, people become both the primary target and the most effective line of defense.

Why Smishing and Vishing Attacks Are Rising

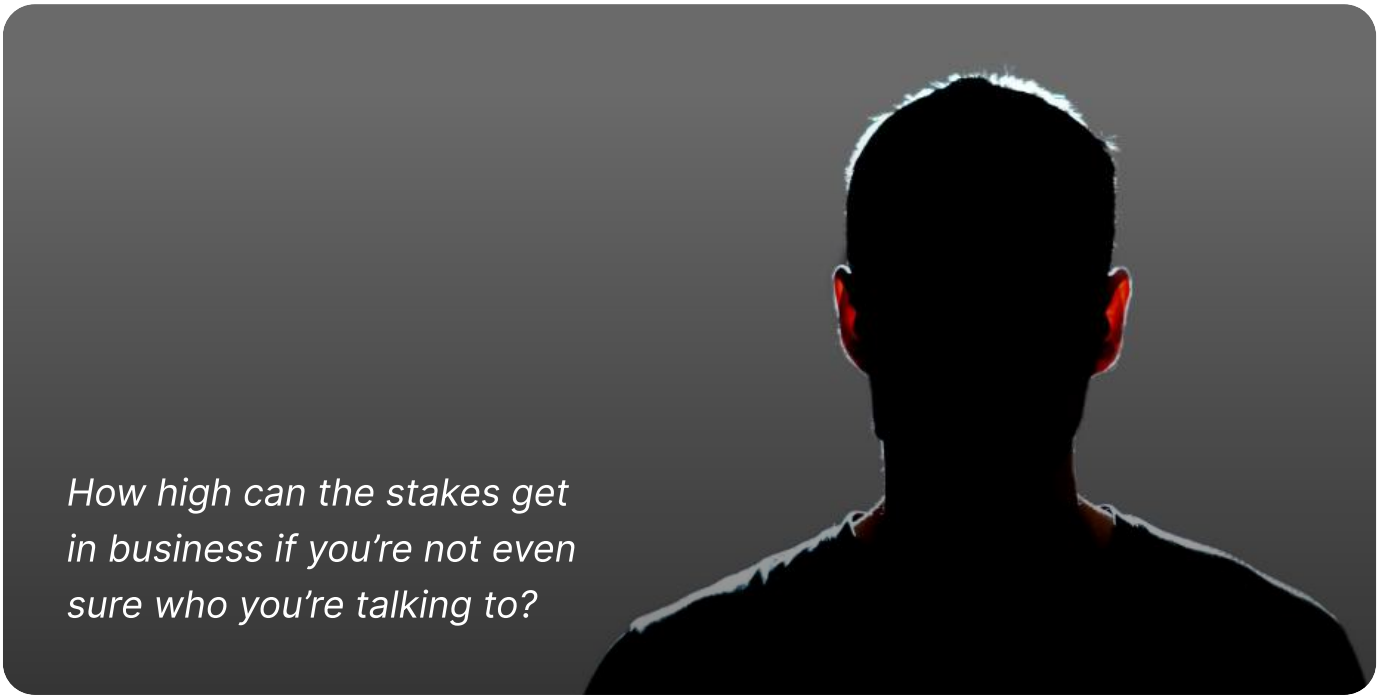

With smishing and vishing, the stakes escalate quickly because the attacker is operating on the device people trust most. Smishing happens over SMS or messaging platforms. Vishing happens over phone calls or callback systems. Both target employee personal devices—often outside the channels security teams monitor most closely.

The stakes are so high because these attacks rely on impersonation, urgency, and misplaced trust.

And the scale is no longer theoretical. Federal reporting shows scam and fraud losses climbing year over year, with the FBI's Internet Crime Complaint Center (IC3) reporting more than $16 billion in losses in 2024, a record that reflects the expansion of social engineering into everyday communications channels. Meanwhile, consumer protection authorities have reported that hundreds of millions of dollars have been lost to scams that began as text messages—$470 million in 2024 alone, according to the FTC.

One example that stands out is a vishing cluster publicly tracked by researchers and federal agencies that used phone-based IT impersonation to compromise Salesforce environments. In these incidents, attackers posed as internal support, convinced employees to grant access or run tooling, and then used the access to exfiltrate data and attempt extortion. It's a clean illustration of the modern pattern: The breach doesn't begin with an email attachment—it begins with a conversation that sounds routine.

Today's threat actors don't improvise. They run playbooks—built from open-source intelligence (OSINT), breached data, and the same workflows your employees follow every day. For individuals, OSINT is what they leave in plain sight: job titles, social profiles, conference appearances, even the way they write. For companies, it's the operational blueprint: org charts, vendor relationships, press releases, support processes, and the details that make business move.

And attackers don't have to guess at phone numbers anymore—they can buy them, scrape them, or pull them from leaks. In November 2024, Amazon confirmed a breach of employee contact data, and reporting tied to the incident indicated the exposed data set contained over 2.8 million lines of employee information, including phone numbers. On a broader scale, the economics are even worse: Wired reported that 3.5 billion phone numbers could be extracted through a WhatsApp contact discovery mechanism, showing how easily phone numbers can be harvested and weaponized at scale.

When that data gets combined with OSINT, impersonation stops sounding suspicious. It starts sounding routine. And most phishing awareness training just can't keep up.

Email is still a major ingress for breaches with a high price, and defense tools on this front have gotten better. Email gateways filter aggressively. Authentication standards help verify senders. Users can more easily slow down and inspect messages.

But SMS and voice don't come with the same guardrails—no authentication layer, no gateway inspection, and very little logging.

There are no equivalents of SPF, DKIM, or DMARC for text messages or phone calls. SMS makes people act quickly without asking questions. It's easy for attackers to raise the stakes and reap the rewards when people have no cues common to an email.

Hundreds of organizations have fallen prey to these tactics, and the price to be paid can run into the hundreds of millions of dollars.

Add generative AI, and the cost to scale drops dramatically. Attackers can personalize at scale—and they only need one person to respond. They no longer have to guess what will sound believable. They can personalize lures using breached data and your company's OSINT trail.

When perpetrators combine that trail with breach data, they can shape messages to match the way real employees:

- Write
- Punctuate
- Ask for help
- Move work forward via text

They can also reference day-to-day activities—ongoing projects, routine approvals, familiar vendors, or internal processes—so the message doesn't just sound authentic. It feels immediately relevant.

That's how one smishing template becomes 10,000 believable messages, each one tuned just enough to look real. Attackers use AI to rewrite the same message thousands of ways, with micro-variations designed to slip past carrier filtering and avoid blocklists. And once a victim answers, what starts as a text can quickly become credential theft, MFA fraud, and account takeover.

Across outreach research and real-world campaigns, text messages routinely generate faster replies than email—which is exactly why attackers favor them.

This is where these attacks win or lose: whether a person responds. SMS interactions feel immediate and personal, and employees are conditioned to treat texts as a higher priority than email. That matters because modern smishing and vishing attacks aren't just about message delivery. They're about engagement. Engagement is what transforms a prompt into a compromise. Across outreach research and real-world campaigns, text messages routinely generate faster replies than email—which is exactly why attackers favor them.

And in real enterprise environments, the gap is measurable. In a Fortune 100 deployment of Adaptive's mobile-first phishing simulations across 12,000 employees, SMS phishing failure was found to be 2x higher than email failure in the first campaign, revealing a blind spot most organizations weren't tracking at all. After remedial training was triggered automatically following failures, the company saw a 54% reduction in SMS phishing failures, and smishing became a permanent part of their security program.

2×

Higher SMS phishing failure, compared to email

54%

Reduction in SMS phishing failures after remedial training

Inside Smishing: How Text Messages Became the New Threat Frontier

IT was requesting a password reset again. The link was short, forgettable, and designed to disappear into the mental clutter of a busy morning. Everything looked authentic—from the helpdesk name to the writing style to the account details.

That is the point of smishing. It's built to look like frictionless internal or third-party communication, arriving in the same channel employees use for real work. And once a message lands on a personal device, the usual defenses don't follow.

In this case, the misdirection was subtle. The URL was masked behind layered redirects and rotating short-link services, making it difficult to see the destination until it was too late.

The attackers could have just as easily used a landing page. Attackers can spin up a convincing replica of an internal login flow in hours (sometimes minutes), complete with branding, phrasing, and the rhythm of real prompts—and they do it at industrial scale. In Q4 2024 alone, the Anti-Phishing Working Group (APWG) recorded 345,881 unique phishing sites in October, 313,288 in November, and 329,954 in December—nearly 1 million unique phishing sites in a single quarter.

Some pages go even further, using device fingerprinting to adjust what the victim sees based on phone type, operating system, or region, so the scam feels native to the device in hand. Threat researchers have also documented how these same phishing kits feed directly into smishing campaigns, enabling high-volume SMS floods through cheap, fast-rotating domain infrastructure and large-scale scam patterns (such as toll-road and payment lures) targeting U.S. recipients.

But even that level of sophistication isn't always necessary. When hesitation appears, attackers escalate quickly—often with a follow-up phone call framed as support "checking in." The goal isn't to trick a filter. It's to keep the target moving, keep the conversation alive, and compress the moment where verification would normally occur. The pressure is the mechanism.

This problem gets worse when attackers don't have to guess who to contact. Personal phone numbers are widely available through data brokers and breached datasets, which makes it easy to target employees outside monitored channels. Regulators have begun cracking down on the sale of personal identifiers like phone numbers, an acknowledgment that this market directly expands the attack surface. For organizations, the implication is simple: if an attacker can find an employee's cell number in minutes, SMS and voice become first-class ingress paths, whether security teams treat them that way or not.

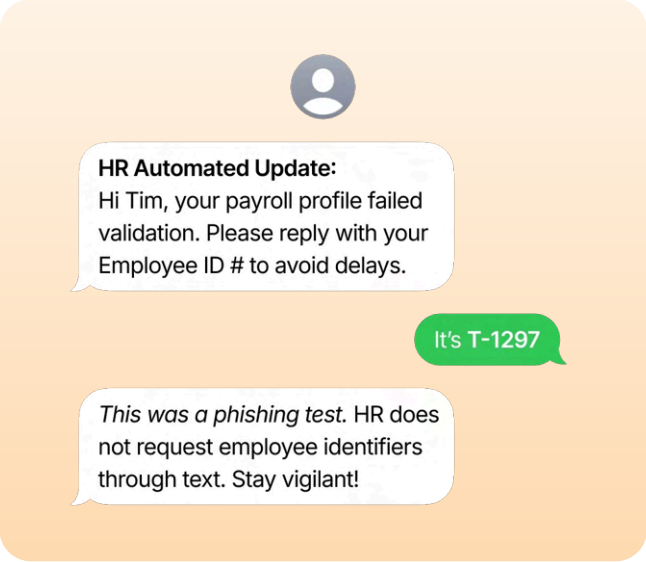

How AI Reinvents the Human Voice for Vishing

In July 2025, U.S. officials disclosed that an attacker used an AI-generated voice to impersonate Secretary of State Marco Rubio, leaving voicemails and messages designed to sound authentic enough to draw high-level targets into conversation. The impostor reportedly contacted three foreign ministers, a U.S. governor, and a member of Congress, using the credibility of a familiar voice and the urgency of official business to probe for access to information or accounts.

That is what makes voice phishing different. When the voice is familiar, people rely on tone, confidence, and recognition more than content, especially under pressure. Most any other human would have done the same.

And the pressure is the point. In these scenarios, attackers exploit the dynamics of live interaction: silence feels uncomfortable, urgency feels personal, and routine requests feel legitimate. This is not a static message you can inspect. It's a conversation designed to keep a person moving forward before doubt can catch up.

Vishing often begins as an escalation. The attacker may start with a text or email, just enough to trigger action, then push the target into a phone call or callback flow where the environment is controlled. Call queues, scripted handoffs, and layered explanations mimic the structure of legitimate workflows.

He didn't question it, because it sounded like work.

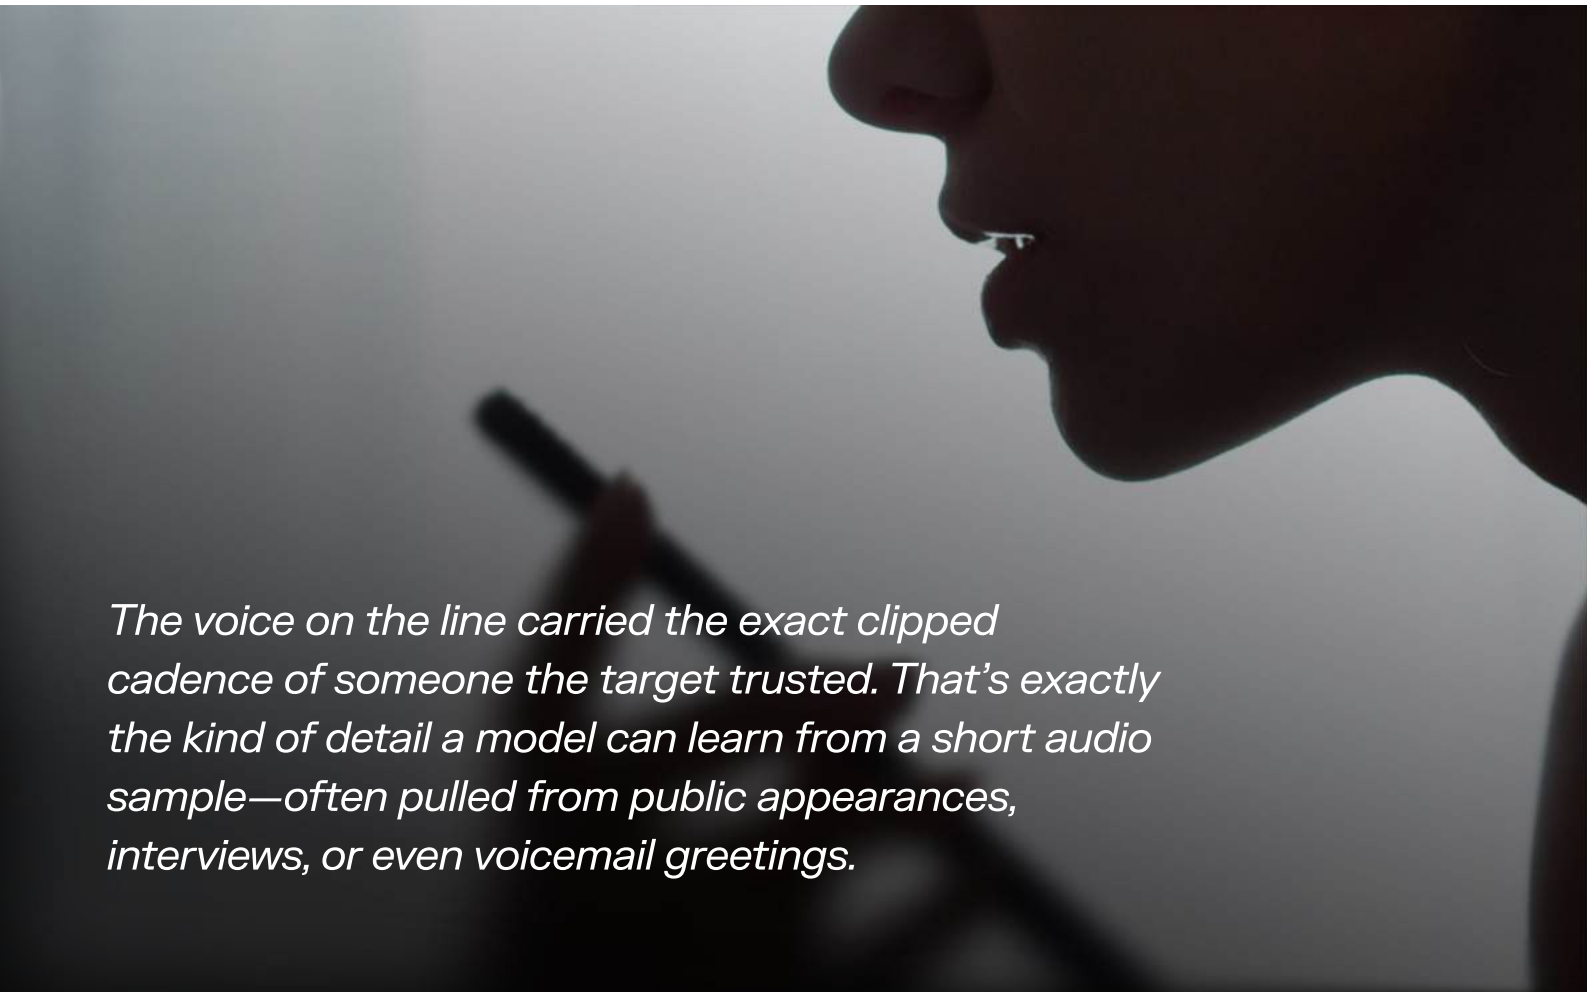

The voice on the line carried the exact clipped cadence of someone the target trusted. That's exactly the kind of detail a model can learn from a short audio sample—often pulled from public appearances, interviews, or even voicemail greetings.

Generative AI has made this threat easier to scale. Synthetic voice models can be trained from short samples pulled from recorded meetings, public videos, or archived audio. While subtle artifacts still exist, they are nearly impossible to detect while navigating a fast-moving request.

In these calls, attackers guide victims through MFA approvals, credential resets, or financial workflows while maintaining the illusion of legitimacy. The call doesn't just deliver the request. It manages doubt.

And that is why incidents like MGM matter. They show how quickly an attacker can turn a support workflow into a breach workflow. The helpdesk isn't just a support function anymore. It's a trust boundary. And voice is now one of the fastest ways to cross it.

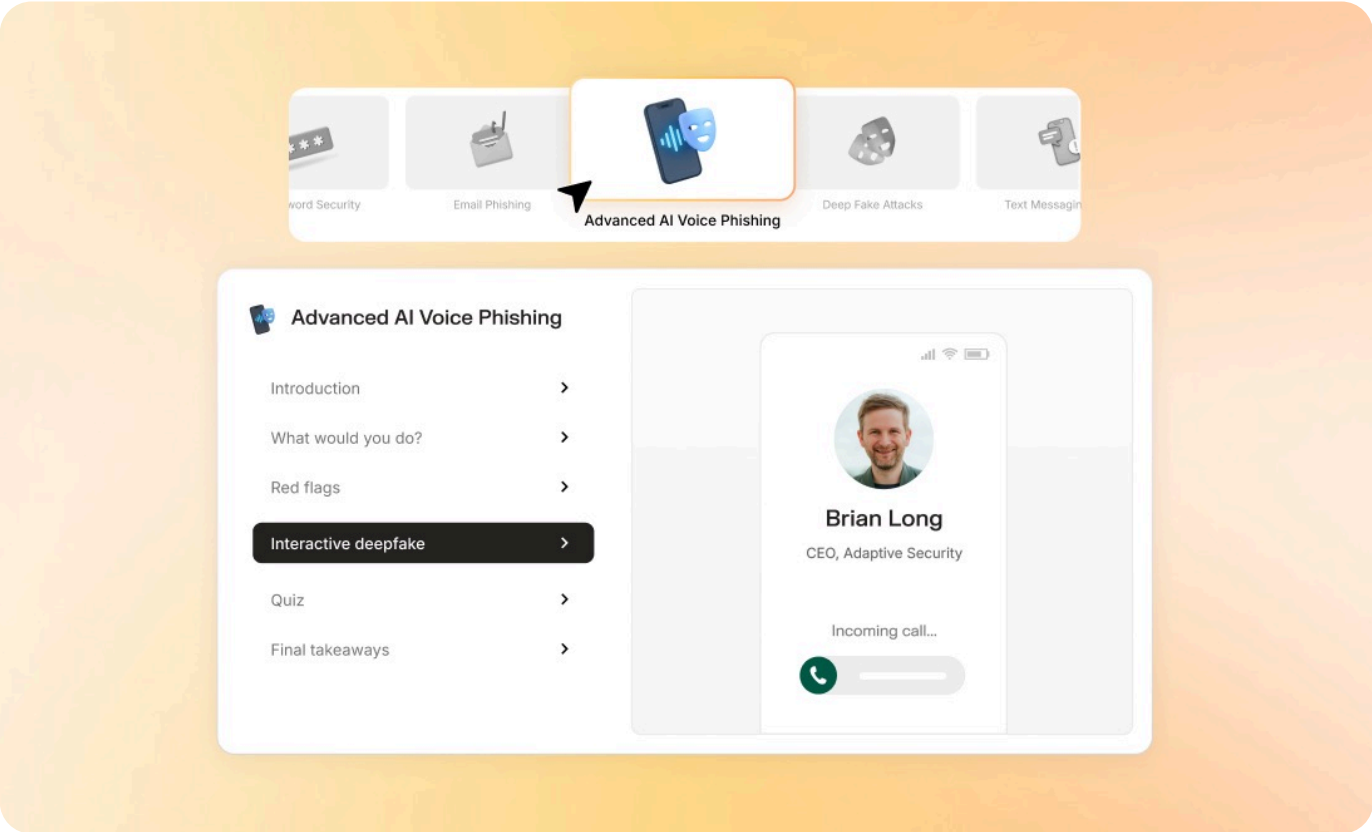

Why Traditional Defenses Fail in the Mobile Era

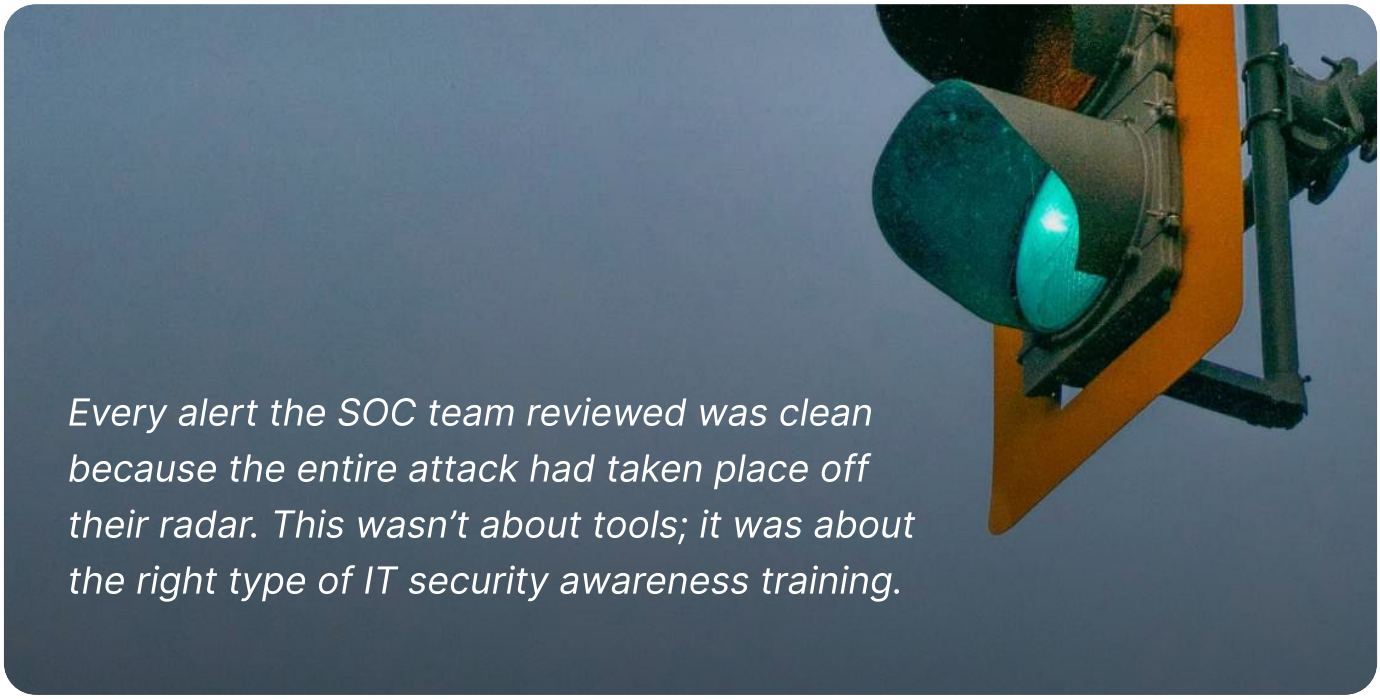

Security teams had built their tooling around email, and for years, that made sense. But filters don't work when there's nothing to filter. SMS and voice channels remain largely unmonitored. National cyber authorities explicitly treat phishing as a cross-channel threat that includes text messages and phone calls, reflecting how attackers increasingly operate outside the inbox. Once a user leaves the inbox, the SOC loses logging, inspection, and context—and the incident becomes invisible until damage shows up elsewhere.

SOC knew there was no reliable inspection layer for SMS. Voice carries no metadata that security gateways can meaningfully analyze. Caller ID spoofing is trivial. SMS protocols lack cryptographic verification. Fraud signaling from carriers is inconsistent and regional.

These gaps are structural, not operational. And personal devices widen them further.

Security teams have little telemetry for inbound call patterns or SMS flows. There's no central logging of mobile interactions. BYOD environments create blind spots that perpetrators deliberately exploit, particularly when employees are off-network and moving quickly.

This is partly historical. Enterprise security evolved around corporate networks, endpoints, and email systems, where organizations could deploy gateways, logging, authentication, and inspection. SMS and voice never received equivalent layers of enterprise-grade verification.

The result is that many organizations treat mobile communications as a "shadow channel"—critical for business, but outside most defensive stacks.

Attackers treat it as the opposite: a direct, low-friction path to the human layer. That shift is reflected across large-scale incident datasets: the Verizon "2024 Data Breach Investigations Report" analyzed 30,458 security incidents and 10,626 confirmed breaches globally, underscoring that human-driven social engineering remains one of the most reliable paths to compromise.

Support impersonation has become one of the most effective vishing techniques because it exploits a workflow designed for speed. A helpdesk agent is trained to resolve access friction quickly, and attackers use that expectation against them—often pairing voice phishing with identity compromise tactics like MFA fatigue, "verify this login" prompts, and device enrollment requests.

This isn't a fringe tactic: In 2024 alone, the FBI's Internet Crime Complaint Center logged 36,002 complaints categorized as Tech Support scams, reflecting how frequently attackers weaponize support workflows to gain access.

More broadly, IC3 recorded 193,407 phishing/spoofing complaints in the same year, reinforcing that impersonation remains the dominant mechanism behind modern social engineering—whether it begins in email, SMS, or voice.

Most security awareness training programs have not adapted to these realities. Annual awareness programs built around email behavior do not map to phone interactions. Mobile interfaces strip away warning cues. Speed replaces scrutiny.

This leaves employees with no realistic exposure to these incidents, forcing them to rely on unrefined instinct—which is exactly what AI phishing is designed to overwhelm.

Creating a Modern, Human-Centered Defense Model

In-depth phishing simulation training had prepared employees to defend against smishing and vishing. The intent of cybersecurity awareness training is not simply to teach caution. It is to build the reflex to pause and verify, especially when a request sounds normal. That preparation cannot rely on static training or email-era assumptions.

These modern information security awareness training programs use AI-driven simulations across SMS and voice. They mirror the attacks employees actually face: the same channels used every day, the same people employees trust, and the same workflows that move work forward. Smishing scenarios reflect real operating rhythms. Vishing simulations recreate pressure without risk. And the key is realism—simulations of people and situations employees actually encounter in daily work.

Over time, employees learn to recognize subtle cues that cannot be reduced to “look for bad grammar.” Through continual exposure, they become better at detecting:

- Abnormal timing
- Mismatched authority
- Unverified callbacks

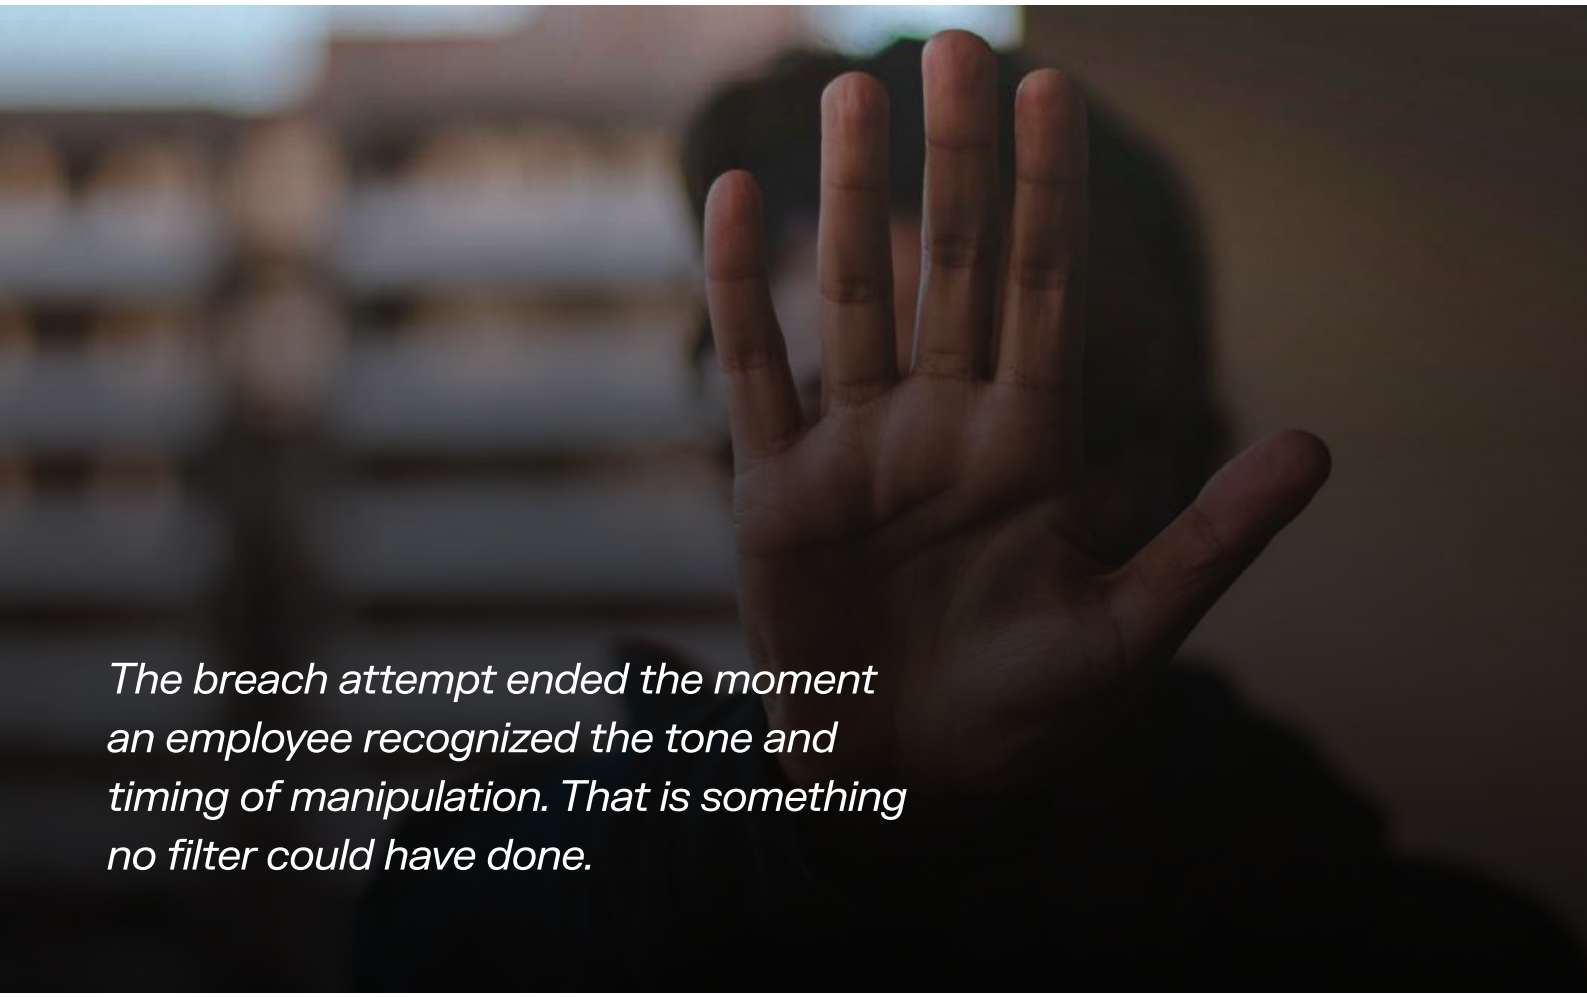

The breach attempt ended the moment an employee recognized the tone and timing of manipulation. That is something no filter could have done.

But the real advantage of modern training is that it can be measured in ways that show whether employees are developing defensive instincts—not just whether they clicked once.

A modern, human-centered defense model should be measured like any other security control: by whether it reduces risk, improves response behavior, and proves progress over time. Core KPIs include:

Human risk score (by role/team)	A clear, role-based view of where the organization is most exposed—so leadership can prioritize high-risk functions like finance, IT, and executives.
Report rate	Whether employees escalate suspicious messages instead of engaging. This becomes early-warning telemetry for security teams and a leading indicator of program maturity.
Time-to-report	How fast employees respond when something feels off. In real-time SMS and voice attacks, speed often determines whether an incident becomes containment or compromise.
Risk delta (improvement over time)	The measurable reduction in human-driven risk over weeks and months—used to demonstrate ROI, track resilience, and prove that training is changing behavior.

Adaptive

Next-Generation Security Awareness

Adaptive Security provides this exact sort of security awareness training, translating frontier AI research into practical, enterprise-safe simulations and behavioral analysis. The platform is designed not only to simulate modern attacks across SMS and voice, but to measure how employees respond—by role, over time, and under real-world pressure. Close alignment with advanced AI research allow the system to strengthen in step with attacker capability, ensuring simulations and training remain credible as threat tactics evolve.

The result is training that elevates the human element into an active defense layer. Employees stop being the weakest link and begin functioning as detection and response sensors, especially in the channels where technical controls are weakest.

Organizations adopting this model report fewer successful smishing and vishing incidents, driven by faster reporting and stronger verification instincts even when the message sounds right.

To learn more about how to protect your workforce against vishing, smishing, and deepfakes, book a demo of Adaptive Security.

Book a demo

Leading organizations trust Adaptive

LENNAR · Figma · ramp · PLAID · nerdwallet · BOSE · stripe · xerox