Adaptive

Phishing Training 2026

How to Spot and Stop AI-Powered Scams

Table of contents

Phishing did not fade away. It learned.

Today's attacks no longer announce themselves with broken English or clumsy impersonation. They arrive as routine work:

- a payroll notice
- a vendor invoice
- a voice message that sounds exactly like a manager asking for a quick approval

AI makes the imitation unnervingly good. The language matches your internal comms. The workflow feels familiar. The sender looks right. The point is not to defeat security software. The point is to move an employee one step past hesitation.

This paper examines how AI-powered phishing has reshaped the threat landscape and why many security awareness programs are struggling to keep pace. Phishing remains a leading entry point, accounting for roughly 40 percent of initial access. It is a reminder that compromise often begins with a human decision, not a technical failure. Employees are asked to make those decisions in ordinary moments, often under time pressure, with fewer reliable cues to lean on.

Most organizations already run awareness training. Where they fall short is relevance. Annual refreshers and generic simulations train people for yesterday’s lures. Current campaigns unfold across email, text, voice, and follow-up interactions that reinforce one another. Security researchers now estimate that more than 80 percent of phishing messages incorporate some form of AI, and delivery has expanded well beyond the inbox. Text messages alone now account for more than a quarter of phishing attempts, and voice-based scams continue to surge.

Defending against this shift means treating training as part of the security system. Employees need repeatable behaviors they can use when uncertainty spikes:

- slow down under pressure
- verify sensitive requests through a second channel
- report suspicious messages early, before the thread spreads

Security teams also need proof these behaviors are taking hold.

Adaptive Security supports this approach by treating phishing defense as a behavioral discipline. Its simulations adapt based on how employees respond, and its measurement focuses on practical risk signals, including:

- which teams and roles are being targeted
- how quickly threats are reported
- whether verification steps are followed
- whether exposure declines over time

In 2026, phishing defense is no longer about teaching people what a scam looks like. It is about preparing them to recognize when reality itself is being manipulated. The organizations that adapt to that truth will be the ones that detect attacks earlier, limit damage, and reduce the likelihood that a single message becomes a headline.

Introduction

It all started when engineers, support staff, and operations personnel at cloud comms platform Twilio began receiving text messages that appeared to come from the company's internal IT systems. The alerts warned of expiring passwords and security changes that required immediate action. The language was precise, the timing plausible. The links embedded in the messages led to pages that closely mirrored Twilio's legitimate authentication workflows.

Nothing about the messages stood out at first glance. In fact, for employees accustomed to frequent security prompts, they blended seamlessly into the background noise of modern enterprise IT.

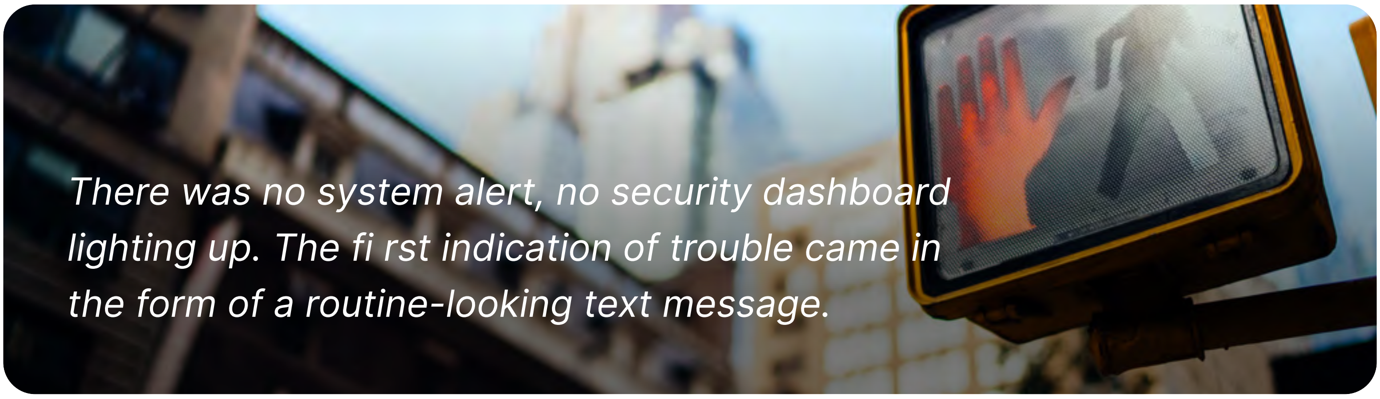

There was no system alert, no security dashboard lighting up. The first indication of trouble came in the form of a routine-looking text message.

Within hours, attackers had obtained valid credentials and gained access to internal systems. Over the following days, Twilio disclosed that attackers had exposed customer data tied to more than 100 downstream organizations. Among those affected were companies like Signal and Okta, both of which also sit at the center of trust-sensitive ecosystems.

The Twilio breach wasn't executed based on an unpatched vulnerability or software flaw. It unfolded through ordinary human behavior exercised under familiar pressure and inside routine workflows. Even in controlled tests, people consistently fail to recognize synthetic media: only 0.1% correctly identified deepfakes, even when warned they were present.

This distinction matters more than ever today.

Because what happened at Twilio was not an aberration. It was a preview of how AI phishing now works when attackers understand that the strongest defenses are no longer technical, but behavioral.

This paper explains how AI phishing has changed the threat landscape and how phishing training has adapted to overcome this evolving risk. It also presents 10 key tips for CISOs to prevent phishing in 2026 and beyond.

When training becomes part of an organization's culture, humans become the most effective defense.

Why Phishing Awareness Training Still Matters in 2026

Many organizations believe their phishing awareness programs are adequate. Completion rates are high. Simulated click rates are tracked. Annual refreshers are documented. From a compliance standpoint, all the boxes are checked.

Real incidents tell a different story.

Employees struggled to question requests that felt familiar or urgent, particularly when those requests arrived through trusted channels or mirrored real workflows. IBM's security research shows that phishing is the most common initial access vector in breaches, accounting for roughly 40% of incidents—underscoring how often compromise begins with human judgment rather than a technical failure. Even experienced teams hesitated when messages aligned closely with how work actually happened.

AI has intensified this gap. Deepfake audio, cloned writing styles, synthetic personas, and AI-generated video now imitate legitimate communication with unsettling accuracy. The old tells that training previously emphasized, the typos, awkward phrasing, or mismatched logos, are no longer reliable indicators.

Instead, deception now hides in subtle behavioral inconsistencies: Timing that feels slightly off, requests that arrive out of sequence, or actions that skip normal verification steps.

AI-powered phishing relies on accuracy, timing, and familiarity delivered across channels. It mirrors how organizations actually communicate, and it only needs to work once. As phishing shifts in this direction, awareness training must function as part of a human-driven perimeter. Security Week's own research team found last year that 82% of all phishing emails they analyzed used some form of AI, a year-over-year jump of 53%.

The uncomfortable reality is that AI-driven phishing has become one of the most efficient paths into the enterprise. It scales cheaply, adapts quickly, and exploits one simple assumption: that internal messages deserve less skepticism than external ones.

Preparing employees to recognize behavioral irregularities is now as essential as strong technical controls — and it works best when paired with them. That's the gap modern security awareness platforms are designed to close. Rather than teaching employees to memorize generic red flags, approaches like Adaptive Security focus on building judgment through role-specific simulations that evolve based on how employees respond and how attacks actually move through the organization.

40%

of security breaches start with phishing

82%

of phishing emails are AI-generated or AI-assisted

The AI Shift: How Attacks Have Changed

AI enables attackers to produce messages that match corporate tone, formatting, and internal jargon. Attachments appear legitimate. Invoices follow real templates. Requests reference actual vendors and systems.

When delivered across channels, these elements blur into a stream of normal business activity.

Organizations can see this change reflected in broader breach data. The human element contributed to 68% of breaches recorded in Verizon's 2024 data breach report. This includes social engineering and credential misuse.

Increasingly, there is no single payload or moment that clearly signals danger. Enterprises no longer face isolated phishing attempts, but orchestrated deception designed to move faster than doubt.

Beyond Email:
Multi-Channel Deception

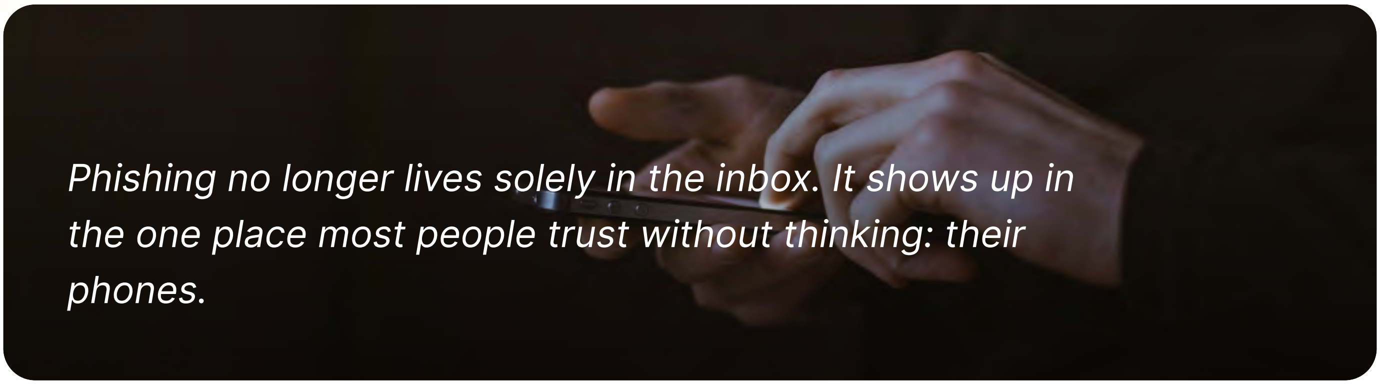

Phishing no longer lives solely in the inbox. It shows up in the one place most people trust without thinking: their phones.

Smishing is climbing fast — more than 20% in recent quarters — and texts now account for over 28% of phishing attacks. Vishing has spiked as well, surging by over 400% as attackers blend voice, SMS, and AI-assisted pressure tactics.

And the damage is real. In 2024, the FTC says consumers reported $470 million in losses to scams that started with text messages — a signal of how profitable fraud outside the inbox has become.

Attackers don't rely on one message anymore. They string together smishing, QR codes, and voice calls to mimic normal workflows, each step making the next one feel easier — and more legitimate.

Security tools still tend to inspect messages in isolation. But AI-assisted phishing doesn't respect those boundaries. When text, voice, and follow-up actions reinforce one another, there may be no single moment that looks obviously suspicious.

This approach exploits trust rather than bypassing controls. It bets that when requests arrive through multiple channels, skepticism drops instead of rising.

Attacks at Cloud Scale

Cloud infrastructure amplifies this effect. Compromised credentials and cloud email services allow attackers to send messages from trusted domains and platforms, blending malicious traffic into everyday operations.

AI speeds up reconnaissance as well. Attackers now use open-source intelligence to gather and map organizational structures, roles, vendor relationships, and work rhythms automatically. They can then tailor campaigns to finance teams, HR, IT, and operations, with messaging aligned to their specific responsibilities.

The result is precision at scale, delivered through infrastructure employees already trust.

The Barrier to Entry for Attackers Has Collapsed

What once required specialized skills is now widely accessible:

- Bad actors can rent infrastructure, reuse credentials, and automate campaigns.
- Off-the-shelf phishing kits replicate enterprise login portals almost perfectly.
- Voice-cloning tools only need seconds of audio to create successful, real-world vishing attacks.

And it's cheap to operate at scale. According to ENISA, modern phishing campaigns rely on low-cost, commodity tooling and "phishing-as-a-service" kits that dramatically reduce both skill and financial barriers for attackers. At the same time, global telemetry from Google and the Anti-Phishing Working Group shows phishing operating at industrial volume, with millions of malicious sites and messages active at any given time. Together, low cost and massive reach have made AI-driven social engineering one of the highest-ROI attack paths into the enterprise.

Technical barrier? There is none.
Attackers simply need patience, context, and automation.

Polymorphic variants are now standard, too. Subject lines, logos, and phrasing can change just enough to evade filters. Every lure appears unique, rendering blocklists ineffective.

What This Means for CISOs

For CISOs, the implication is clear. When visual and textual cues disappear, they must train employees to use their intuition. Phishing awareness must reflect how attacks unfold across channels, under pressure, and without obvious errors—because phishing (including vishing, malspam, and malvertising) accounts for about 60% of observed intrusion vectors in ENISA's recent threat landscape analysis.

This requires a different approach to simulation, measurement, and culture. Training must evolve alongside threat behavior, and security teams must track employee progress in ways that reflect actual risk, not theoretical success.

That shift begins with practical guidance employees can rely on when uncertainty pops up.

CISOs must adapt to the reality that preparing employees to recognize behavioral irregularities is now as critical as any technical control.

Effective training going forward is all about prevention.

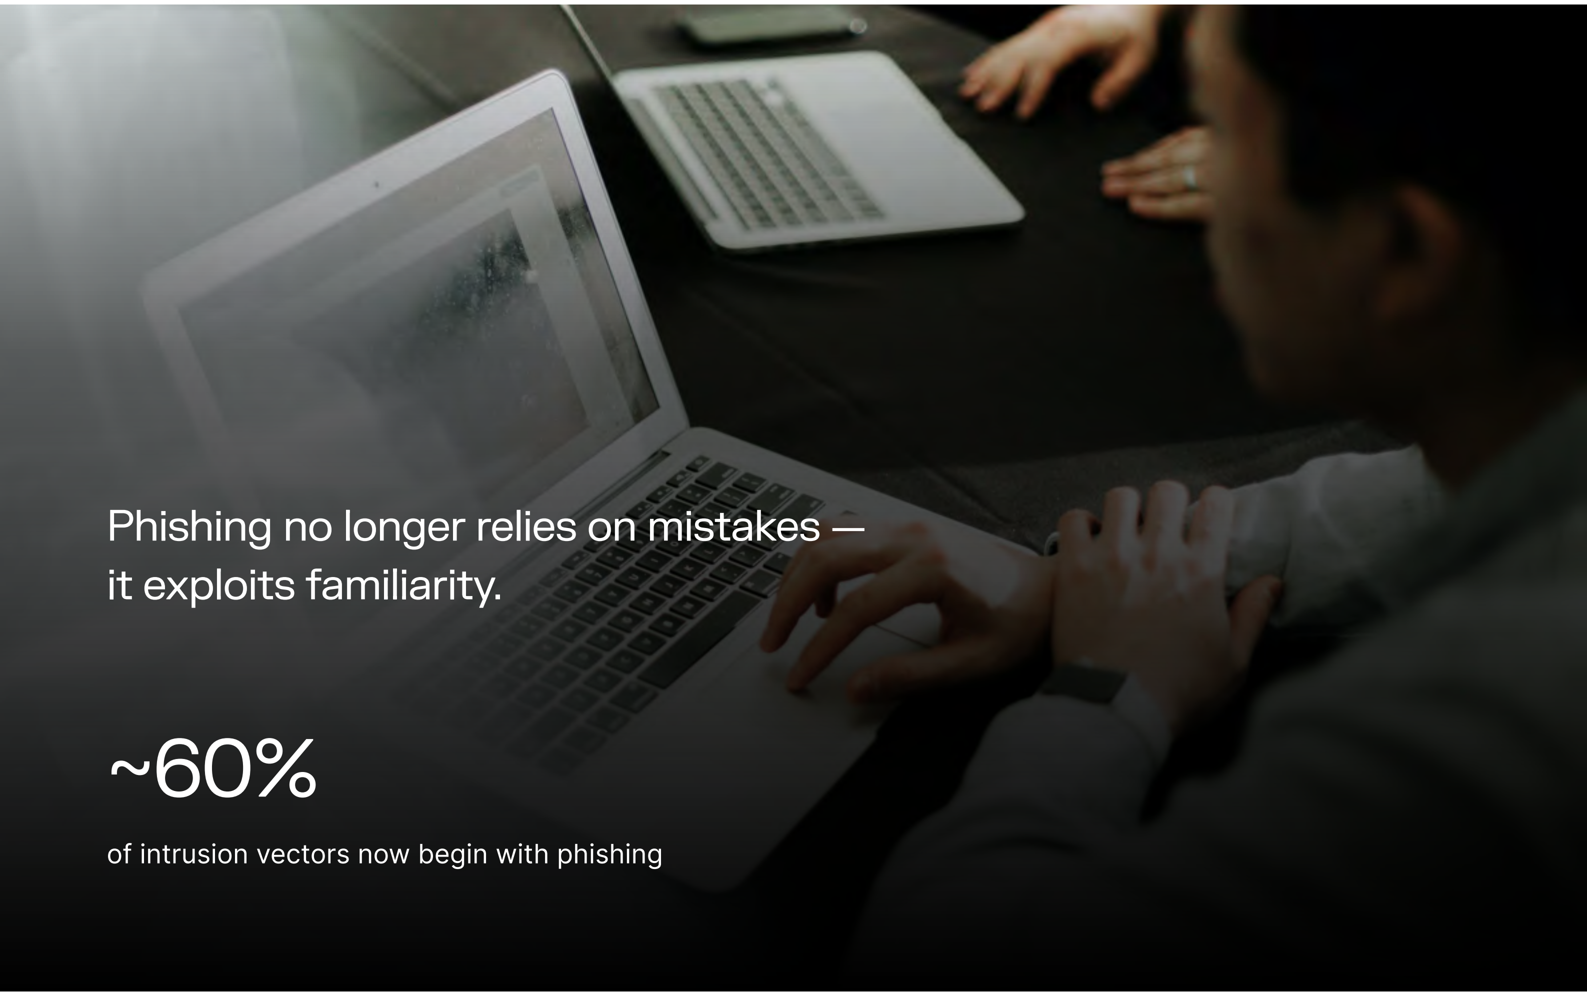

Phishing no longer relies on mistakes — it exploits familiarity.

~60%

of intrusion vectors now begin with phishing

10 Must-Know Tips to Prevent Phishing in 2026

CISOs can start their journey to effective security training with the following employee guidelines: 10 brief tips that will hold up under pressure, across email, text, voice, and video.

These will help transform incidents into guidance without overwhelming employees with theory.

1. Verify Unexpected Requests, Even from Familiar Sources

Familiar names and internal tone are no longer proof. Confirm both the request and the result through a second, official channel before acting.

2. Hover Before Clicking and Inspect Destinations

Links often mask their true destination behind redirects or shorteners. A brief pause can reveal a mismatched domain.

3. Slow Down When Urgency Appears

Urgency is a common manipulation tactic. When pressured to act quickly, verify.

4. Validate Identity Before Responding to Voice or Video Instructions

Convincing voices and videos can be synthetic. Treat unexpected instructions as unverified until confirmed elsewhere.

5. Check for Context Mismatches

Many scams are *almost right*. Timing, workflow, or content that feels slightly off often signals manipulation.

6. Report Suspicious Messages Immediately

Early reporting helps security teams spot broader campaigns and limit impact, even when they take no action.

7. Use Multi-Factor Authentication Consistently

MFA cannot stop every attack, but it often prevents stolen credentials from becoming a full-blown compromise.

8. Use AI-Powered Security Tools That Detect Behavioral Anomalies

Modern tools look beyond content, correlating behavior and context to reveal subtle warning signs.

9. Embrace Continuous Security Awareness Training, Not Annual Refreshers

Short, recurring phishing training exercises reinforce instincts and keep pace with changing attack methods. The best programs add on-the-spot coaching tied to the exact mistake or hesitation the employee just made.

10. Don’t Be Afraid to Question

Employees should feel supported when they pause or challenge unusual requests, regardless of seniority.

Taken together, these practices shift phishing defense from a checklist to a mindset, one that favors hesitation over speed and verification over assumption.

The Cost of Getting It Wrong

When AI-driven phishing escalates, the damage compounds quickly. This can mean financial, operational, and regulatory costs that are intertwined in a modern enterprise.

Financial Consequences

Phishing accounted for hundreds of thousands of complaints and billions of dollars in losses, according to the FBI's 2024 IC3 report. This shows how quickly small decisions translate into major effects in material ways.

Operational Disruption

The operational disruption of AI phishing, regardless of how successful it is, still has a financial impact. Containment may mean freezing accounts or taking systems offline. This not only slows internal operations and processes but also service and product delivery.

Businesses lose more time and money long after the original phishing message is gone.

Regulatory and Legal Exposure

Legal and compliance teams get involved once the possibility of data exposure surfaces. If AI phishing results in any type of data loss involving PII, regulatory organizations call for immediate action and fines.

Case in point: Britain's Information Commissioner's Office (ICO) fined U.K.-based Interserve £4.4 million ($6 million) after an employee downloaded a phishing email. The incident, according to The Guardian, left 283 systems and 16 accounts compromised.

Hackers were able to snatch the personal and financial data belonging to as many as 113,000 employees. Decisions about disclosure, scope, and timing carried weight as a security issue became a governance issue.

Reputational Harm

Operational disruptions and regulatory corrections ultimately become public-facing problems. These shocks hit the bottom line in measurable ways.

These shocks hit the bottom line in measurable ways. After disclosure of a major breach at Coupang, the e-commerce company saw roughly $8 billion wiped from its market capitalization amid investor concerns about cyber risk and regulatory fallout.

Research also shows that publicly traded firms commonly experience share-price underperformance in the weeks and months following breach announcements, reflecting persistent investor anxiety about reputational and operational damage.

The same happens to third-party partners that fear disruptions and financial fallout from reputational hits to the primary enterprise. In the long run, everyone's bottom line and market share are at stake.

Why Enterprise Maturity Does Not Guarantee Safety

> *The maturity of ingress and egress security protection systems and filtering has very little bearing on effectively stopping AI phishing.*

The same can be said for "mature" training that uses outdated simulations as part of a standard training program every four months. In the face of AI-driven phishing, this approach no longer works. And by the time the cost is clear, the moment to stop it has already passed.

Adaptive Security as Your Training Partner

The problem for most enterprises is that security awareness training for employees can't keep up with threats. Most simulations are abstract, sanitized, and predictable. But actual attacks are fast, contextual, and personal. They move across email, text, voice, and video without announcing themselves as malicious.

Effective training requires simulations that mirror modern phishing as it actually happens: multi-channel, role-aware, and designed to build judgment under pressure. A finance employee faces different lures than an IT administrator. An executive gets different requests than a frontline worker. Training has to reflect those realities.

Adaptive Security delivers simulation-led training built around that standard. Scenarios evolve based on how employees respond, and the program measures what matters in practice—like reporting behavior, speed, and decision patterns—not just clicks.

And when employees report suspicious messages, Adaptive can help close the loop with automated phish triage—so reporting becomes a faster path to validation and response. Adaptive's platform aligns with rontier AI research to keep simulations current without turning training into theory. The goal is practical parity: if attackers are using AI to refine deception, defenders need training that reflects the same sophistication.

Recommendations for Leaders

Organizations should begin by evaluating whether their training programs merely respond to phishing incidents or actively anticipate them. That distinction matters. Preparing for modern phishing requires a shift not only in mindset, but in how readiness is measured.

Modern Security Training

Security awareness training must operate as part of a human-driven perimeter, not as a compliance exercise. Today's attacks unfold across text, voice, QR codes, and follow-on interactions. Simulations that focus only on email no longer reflect the conditions employees face.

Effective programs expose employees to these patterns deliberately and repeatedly. Short, recurring exercises help employees practice hesitation, verification, and escalation in real time, rather than relying on memory from annual refreshers. Over time, this repetition builds familiarity with uncertainty itself, which is what attackers exploit.

Measuring What Actually Reduces Risk

Click rates alone cannot explain whether an organization is becoming more resilient. What matters is whether employees surface suspicious activity, how quickly they do so, and whether high-risk requests are verified before action is taken.

Adaptive Security was built to support this shift by tracking behaviors that correlate directly with breach prevention. Key metrics include:

- **Human Risk Score:** Identifies which teams and roles are most targeted and most vulnerable, so risk is visible where it concentrates.
- **Report Rate:** Measures whether employees surface suspicious messages instead of silently deleting them, improving detection and response.
- **Time-to-Report:** Tracks how quickly potential phishing reaches security teams, reducing dwell time and limiting spread.
- **Verification Compliance Rate:** Measures whether employees follow pause-and-verify workflows for sensitive payment, access, or data requests.
- **Risk Delta:** Shows whether risk is actually trending down over time, quarter after quarter, rather than staying flat.

Together, these signals allow security teams to move from anecdotal assurance to measurable improvement.

Adaptive

Building a Culture That Supports Verification

Metrics alone are not enough. Employees must feel supported when they pause, verify, or escalate something that feels off. When reporting is treated as a contribution rather than a disruption, organizations see faster detection and fewer downstream impacts.

The Adaptive Security platform helps reinforce this culture by pairing measurement with feedback and adjustment, resulting in fewer escalations and more people speaking up early. Over time, verification becomes routine, reporting becomes timely, and security awareness becomes embedded in how work gets done. With Adaptive, employees will be prepared to recognize scams, even as they evolve.

This is what phishing defense looks like in 2026. Not a once-a-year module, but a living system that is observed, tested, and improved continuously.

To stop AI-powered cyber attacks, learn how Adaptive Security uses multi-channel phishing simulations and behavior-based metrics to reduce risk over time.

Click here

Leading organizations trust Adaptive

LENNAR Figma ramp PLAID

nerdwallet BOSE stripe xerox