

Kaseya®

2026 REPORT

Email security report: AI, phishing and the new era of defense



Table of contents

Email security at an inflection point	3
The state of cybercrime: Escalation, evolution and email's central role	4
Brand impersonation: The trust layer under attack	9
Threat evolution: New categories, new behaviors	14
The structural limits of legacy email security	21
AI has changed the shape of phishing	22
Predictions for 2026: Escalation, consolidation and AI at scale	25
Closing outlook	28

Email security at an inflection point

Cybercrime has become a structural feature of the digital economy.

Phishing continues to be the most prevalent and impactful attack vector. Early 2025 telemetry from security research groups shows a marked rise in AI-enabled campaigns and infrastructure-based impersonation, indicating that the threat landscape is still in flux. But the story is not just about growth. It's about transformation.

Generative AI has reshaped phishing from a high-volume nuisance into a precision instrument. Messages are polished, contextual and increasingly delivered through trusted infrastructure. Traditional detection signals — poor grammar, suspicious domains, obvious malicious links — are now unreliable. In their place: intent-driven manipulation embedded in legitimate workflows.

At the same time, infrastructure risk is concentrating. Major cloud and SaaS platforms now underpin global commerce and communication. The compromise of a single ecosystem can have cascading effects far beyond any individual organization.

This report examines how phishing has evolved, where AI has accelerated both offense and defense and what organizations should anticipate in 2026.

The conclusion is not that email is broken. It is that email security has entered a new phase — one defined by adaptive detection, contextual reasoning and systemic resilience.



5 key findings

- 1 Phishing is now AI-generated by default
- 2 Brand impersonation is the dominant fraud lever
- 3 Infrastructure abuse replaces malicious domains
- 4 Attack intent matters more than traditional indicators
- 5 Static detection models are obsolete

The state of cybercrime: Escalation, evolution and email's central role

Cybercrime is no longer a rising threat — it's a sustained and accelerating economic force. According to the FBI's 2024 Internet Crime Report¹, total reported losses from cybercrime reached \$16.6 billion, representing a staggering 295% increase over the past five years.

This dramatic growth shows not only increased attack volume, but also the professionalization and industrialization of cybercriminal operations. Today's attackers operate at scale, use automation and AI and increasingly target the most efficient attack surface available: email.

Phishing remains the primary attack vector

Phishing continues to be the leading method cybercriminals use to gain access to organizations and people. In 2024:

26% of all cybercrime complaints filed with the FBI were phishing-related

274% increase in phishing losses
\$18.7B → \$70B year over year

These numbers underscore a fundamental reality: email is still the most scalable and cost-effective attack channel available to threat actors. Rather than relying solely on malware or infrastructure exploits, attackers increasingly exploit trust, urgency and human decision-making.

While some categories of cybercrime showed decline — including a 79% decrease in ransomware-related losses and a 34% reduction in reported data breach losses — this doesn't indicate that ransomware or data theft are disappearing. Instead, it reflects a strategic shift: attackers are increasingly applying phishing and Business Email Compromise (BEC) schemes as lower-risk, high-return alternatives to disruptive encryption-based attacks.

¹At the time of publication, the FBI's 2025 Internet Crime Report has not yet been released. The 2024 report therefore provides the most recent comprehensive federal dataset available for trend analysis.

Cybercrime by the numbers

\$16.6B
total losses

295%
five-year growth

26%
phishing complaints

274%
phishing loss growth

Business email compromise: High impact, high precision

Among phishing-related threats, BEC remains one of the most financially devastating.

- \$2.8 billion in reported BEC losses
- \$129,193 average loss per incident

Unlike broad phishing campaigns, BEC attacks are highly targeted. Threat actors often impersonate executives, vendors or partners — often fueled by prior account compromise or conversation monitoring. The high average loss reflects the strategic nature of these attacks, which often target financial workflows such as wire transfers, invoice payments and payroll changes.

As explored in previous sections of this report, account takeover (ATO) often serves as the driving mechanism behind BEC, allowing attackers to weaponize legitimate email accounts and bypass traditional reputation-based defenses.





The expanding fraud ecosystem

While phishing dominates the cybercrime landscape, other financially motivated fraud schemes are also accelerating.

- Investment fraud involving cryptocurrency rose 145%, reaching \$9.32 billion
- Call center scams accounted for \$1.9 billion
- Gold courier scams totaled \$219 million
- Emergency scams accounted for \$2.7 million

These figures reflect a broader shift toward social engineering-driven fraud — scams that exploit urgency, fear and trust rather than technical vulnerabilities. Many of these schemes begin or are supported by phishing-based email campaigns.

Where the money is going

Phishing losses

\$70 Billion

Crypto investment fraud

\$9.3 Billion

BEC losses

\$2.8 Billion

Call center scams

\$1.9 Billion

Ransomware

\$59.6 Million


Ransomware: Changing tactics, not disappearing

Although reported ransomware losses declined by 79%, ransomware remains a persistent threat — particularly to critical infrastructure sectors.

The top five ransomware variants



The top infrastructure sectors impacted



The reduction in losses may reflect improved backup strategies and response readiness. However, ransomware actors increasingly combine data theft, extortion and phishing-based initial access techniques — reinforcing email's role as the gateway to broader compromise.


Small and mid-sized organizations at high risk

While large enterprises often dominate headlines, smaller organizations face disproportionate impact.

For many SMBs, a single BEC or ransomware event can threaten operational survival. Limited security staff, legacy email infrastructure and reliance on third-party vendors also increase exposure.

This makes modern email security — particularly advanced phishing and impersonation detection — a critical control for organizations of all sizes.

The data is clear. Phishing is no longer one category of cybercrime; it's the primary engine driving financial loss across industries. As attackers adopt AI, refine impersonation tactics and target trusted workflows, organizations should rethink how email threats are detected and stopped. The following sections examine how brand impersonation, emerging threat categories and AI-driven attacks are reshaping the email security landscape.



of ransomware attacks targeted organizations with fewer than 1,000 employees

SMB median BEC loss:

\$50,000

60%

of small businesses close within six months of a cyberattack

Brand impersonation: The trust layer under attack

If phishing is the leading cybercrime category, brand impersonation is its most effective delivery mechanism.

The top 25 most-phished brands of 2025 represent the organizations attackers most frequently impersonated in email-based campaigns detected by INKY, Kaseya's advanced email security solution. These brands are not random — they are high-trust institutions embedded in daily business operations, financial transactions and personal workflows.

The scale of brand-based phishing in 2025

The numbers reveal the magnitude of the problem:

- 4.5+ billion emails processed in 2025
- 6,688,601 brand impersonation emails detected in H2 2025 alone
- 281 unique brands impersonated
- 5,283,200 phishing emails associated with the top 25 brands

While these 25 brands represent only a subset of the 281 organizations impersonated, they account for a disproportionate share of brand-based phishing activity. All of these attacks were successfully detected and blocked before reaching end users, preventing downstream impacts such as credential theft, financial fraud, business email compromise and ransomware intrusion.



Brand impersonation at scale




Fresh Phish Spotlight

Cross-brand impersonation

In a Zoom-themed phishing campaign, attackers combined trusted corporate branding with government authority, impersonating Social Security Administration alerts delivered through Zoom infrastructure.

This blending of trusted brands and official institutions reflects a growing tactic: multi-layered impersonation designed to bypass skepticism.



What's changed since our last report?

Brand impersonation is not new. What's changed is the quality, realism and contextual sophistication of these attacks.

Three major shifts stand out:

1 Visual fidelity increased dramatically

AI-generated layouts, brand-consistent typography and pixel-perfect logos now closely mirror legitimate communications. Traditional detection methods that rely on sender reputation or URL analysis struggle to identify these emails when infrastructure appears legitimate.

2 No-payload phishing became common

Increasingly, brand impersonation emails omit traditional malicious elements such as links or attachments. Instead, they:

- Provide phone numbers (voice phishing pivot)
- Request replies
- Use QR codes
- Direct users to "verify" via secondary channels

These techniques reduce detectable indicators while increasing reliance on user decision-making.

3 Brand impersonation expanded across sectors


While financial and technology brands remain heavily targeted, 2025 showed continued expansion into:

- Retail (shipment notifications)
- Telecommunications (account alerts)
- Logistics (delivery failures)
- Professional services platforms

Attackers increasingly impersonate brands that align with everyday workflows, making phishing attempts blend seamlessly into normal business operations.



Top impersonated brands

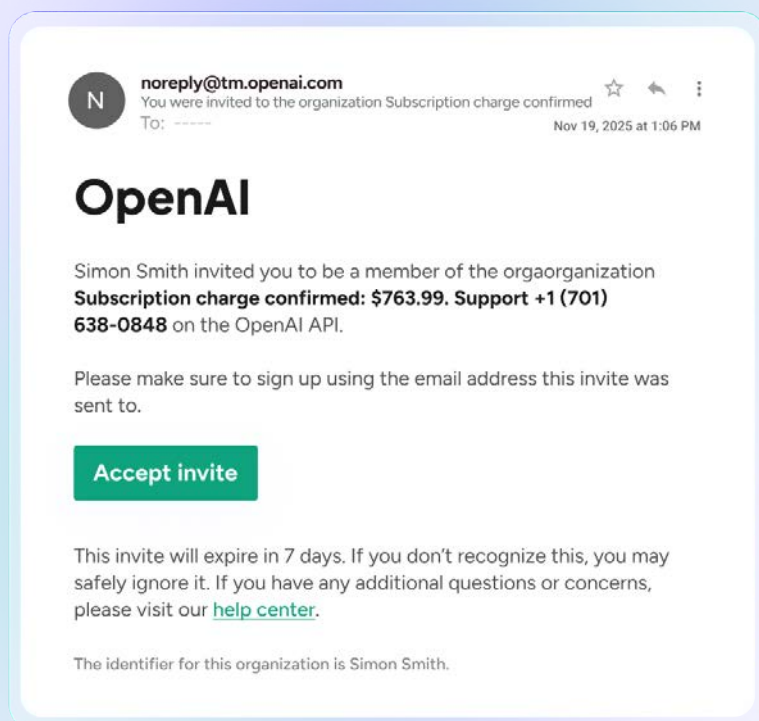


Fresh Phish Spotlight

Platform abuse meets callback phishing

Threat actors used SendGrid — a legitimate email delivery service — to distribute OpenAI-themed phishing messages. Rather than including malicious links, the emails directed recipients to call a phone number, shifting the attack into voice-based social engineering.

This reflects a broader 2025 pattern: attackers increasingly combine trusted infrastructure with callback phishing to evade URL-based detection entirely.



Why brand impersonation works

Brand impersonation succeeds because it transfers trust. Recipients don't evaluate an email in isolation. They interpret it through:

- Recognition of the brand
- Familiar workflows
- Perceived urgency
- Contextual relevance

When attackers mirror these cues convincingly, traditional email defenses that focus on domain reputation or known threat lists become insufficient.

This is particularly critical in BEC and account takeover scenarios, where impersonation of high-trust brands often precedes credential theft or financial fraud.

Brand impersonation is not merely a phishing tactic — it's the psychological lever that drives most financially motivated email attacks.

Brand impersonation is only part of the story. In 2025, entirely new threat categories emerged alongside increasingly sophisticated phishing techniques. The next section examines the new threat categories that surfaced and how attackers are adapting their methods to evade legacy detection models.

Threat evolution: New categories, new behaviors

While brand impersonation stayed dominant, our 2025 data showed meaningful shifts in attacker behavior that required new detection models and new threat classifications.

These shifts were not incremental. They reflect a fundamental evolution in how phishing campaigns are created, delivered and scaled.

New threat categories

To address emerging attack techniques, INKY introduced several new detection categories:

Internal name match

Attackers increasingly spoofed internal display names to mimic legitimate employees. Even when the sending domain differed, matching the internal employee's name created a powerful illusion of authenticity.

This tactic is especially effective in distributed and hybrid workforces, where employees may not personally know every colleague.

Automatically released from quarantine

Attackers began designing emails specifically to bypass or trigger automatic release workflows in legacy systems. Messages engineered to appear benign at first pass were released post-analysis, exposing users to delayed threats.

This category reflects attacker awareness of security tooling behavior — not just user behavior.

Abused service

One of the most notable shifts in 2025 was the abuse of legitimate platforms and trusted infrastructure. Attackers leveraged:

- Microsoft Teams
- PowerBI
- Cloud-hosted document platforms
- Online form services
- Shared collaboration tools

Rather than hosting phishing content on obviously malicious domains, attackers embedded malicious intent within trusted services — often inserting phone numbers, alternate contact instructions or secondary payload links in overlooked fields.

This shift represents a strategic move away from suspicious infrastructure and toward trusted ecosystems.


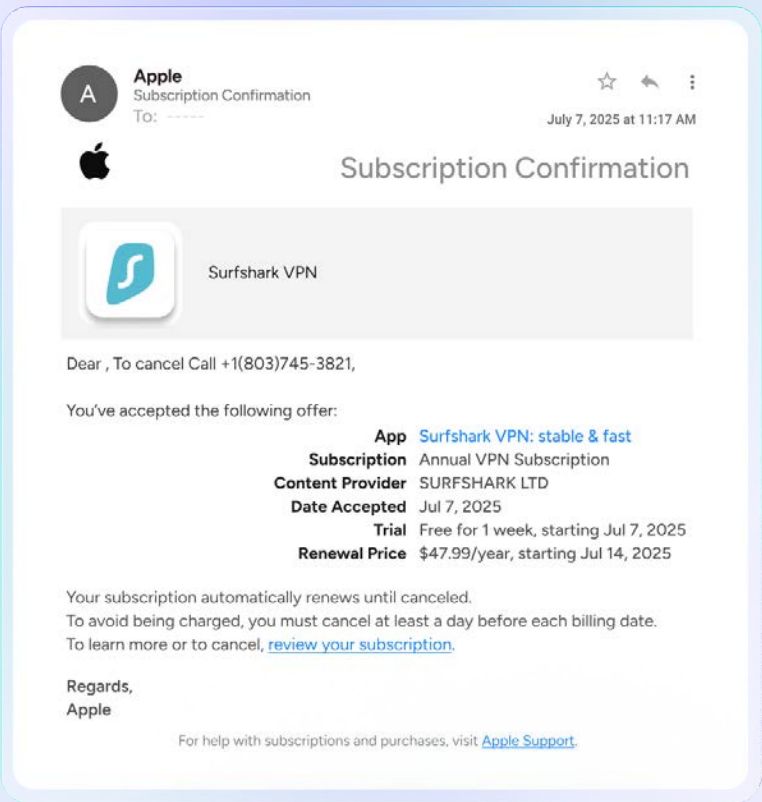
Fresh Phish Spotlight

When authentication isn't enough

In one campaign, attackers used DKIM replay techniques to resend legitimate emails while modifying contextual elements. Because the original message was properly authenticated, traditional filters viewed the replayed message as trustworthy.

Similarly, invoice abuse campaigns impersonating Apple and PayPal used legitimate billing systems to generate authentic-looking payment notifications. The emails contained no malicious attachments and passed standard authentication checks — yet the intent was fraudulent.

These examples illustrate a key shift: attackers are increasingly exploiting trust signals rather than attempting to bypass them.



What dominated in volume:
Threat telemetry

While new techniques emerged, telemetry shows which threat classes dominated overall volume.

Top red-danger categories in 2025
(raw counts converted to % for visual clarity):

- Phishing content
- Spoofed internal sender
- Brand impersonation
- Potential sender forgery
- Suspicious URL
- Reported phish

Notably, phishing content alone accounted for more than 31 million detections, reinforcing email as the primary delivery channel.


However, the real insight is not just the volume — it’s the layering.

Many attacks triggered multiple labels simultaneously:

- Brand impersonation + suspicious URL
- Internal spoof + sender forgery
- Phishing Content + Abused Service

This reflects a growing trend: attackers are stacking techniques within single campaigns.

Primary email threat categories observed in 2025



The stand-out takeaway: AI-driven variability

One of the most significant 2025 observations from INKY telemetry: attackers are no longer running single-template campaigns. Historically, phishing kits were static. A campaign might use one template repeatedly, creating detectable repetition patterns.

In 2025, AI changed that. Attackers used AI to generate:

- Invoice phishing
- Voicemail phishing
- Fax-themed phishing
- Payment request phishing

—all originating from the same sending domain and often using the same malicious link, but with completely different messaging and layouts.

This variability made campaign-based detection far more difficult. Detection signals shifted from surface indicators to underlying intent.

INKY responded by expanding sender profiling and intent-based detection — analyzing patterns across domains and message behavior rather than relying on content repetition.




Fresh Phish Spotlight

Contextual phishing in real time

In early 2025, attackers launched DHS impersonation campaigns aligned with current U.S. deportation news cycles. The emails were highly contextual, professionally formatted and emotionally charged.

This type of real-time event-driven phishing demonstrates how AI allows attackers to rapidly generate context-aware campaigns that align with breaking news.



Additional emerging patterns in 2025

Abuse of trusted infrastructure

Attackers increasingly leveraged legitimate Microsoft ecosystem services and collaboration platforms. The abuse of trusted infrastructure reduces traditional red flags and increases user trust.

Protected and encrypted payloads

Attackers increasingly delivered:

- Password-protected PDFs
- Encrypted attachments
- Files requiring secondary authentication

These files often contained malicious links or instructions once opened.

INKY detects these by:

- Safely rendering content
- Extracting embedded URLs
- Evaluating contextual intent
- Applying computer vision to protected document previews

Calendar phishing

A growing tactic involves embedding malicious links within calendar (ICS) attachments.

Because many email clients automatically add calendar invitations to users' schedules, malicious events can appear without the recipient explicitly clicking the email.

Users often interact with the calendar entry later, triggering phishing payloads.

INKY detects these by analyzing:

- Embedded URLs inside calendar attachments
- Behavioral anomalies
- Intent signals

This category is expected to grow significantly in 2026.





Phishing in 2025: From volume to precision

The force driving this transformation was artificial intelligence. The defining shift in 2025 was not simply increased phishing volume, it was increased sophistication.

Attackers moved decisively toward:

- Trusted infrastructure abuse
- Multi-channel pivoting
- AI-driven message variability
- Intent-focused manipulation

Phishing campaigns became more precise, less repetitive and more context-aware. Emails authenticated properly. Infrastructure appeared legitimate. Language was polished and error-free.

Rather than defeating security controls, attackers increasingly used the trust signals those controls were designed to protect.

Email remained the primary entry point, but it was no longer the entire attack surface. Calendar invitations, collaboration platforms, protected documents and callback phone numbers extended phishing beyond the inbox.

The net result: traditional detection approaches built to flag obvious indicators struggled against attacks that appeared technically legitimate but were strategically malicious.

The structural limits of legacy email security

The evolution of phishing in 2025 exposed structural limits in traditional email security solutions.

Legacy detection systems were designed to identify technical anomalies — suspicious domains, malicious URLs, poor grammar and repeated templates. For years, those signals were reliable because phishing campaigns relied on scale, automation and visible imperfections.

That operating model no longer reflects the threat landscape.

As attackers increasingly authenticate their infrastructure, abuse trusted SaaS platforms and generate AI-polished content at scale, historical indicators of compromise have become less predictive. Messages pass SPF, DKIM and DMARC checks. Links resolve to legitimate cloud services. Language is grammatically correct and contextually appropriate. Campaigns no longer rely on identical templates but instead use AI-driven variability to evade clustering.

The gap is not the absence of security controls. It’s a mismatch between detection models built for a previous generation of threats and the intent-driven manipulation that now defines modern phishing.

Why legacy email security models fail

Legacy Approach

Now

Reality URL reputation



Trusted infrastructure abuse

Grammar detection



AI-polished content

Template clustering



AI-driven variability

Domain blacklists



Infrastructure reuse

Static rule engines



Intent-driven attacks

As phishing becomes **less about technical compromise and more about behavioral manipulation**, detection must move beyond surface indicators and toward contextual reasoning.

AI has changed the shape of phishing

Artificial intelligence is the force driving the structural shift in phishing.

In 2025, generative AI permanently altered the economics of email-based attacks. The cost of producing convincing phishing content collapsed. What once required skill, time and manual iteration can now be generated instantly.



Attackers can produce

- Grammatically flawless messages
- Brand-consistent formatting
- Multiple campaign variants from a single prompt
- Context-aware messages tied to current events
- Personalized communications aligned to specific roles or workflows

Industry research indicates

- 83% of phishing emails now contain AI-generated or AI-assisted content²
- 40% of BEC emails are generated using AI³
- AI-generated phishing emails achieved a 54% click rate, compared to 12% for generic campaigns.⁴

The impact is not incremental. AI eliminated the historical tradeoff between scale and realism. **Attackers now have both.**

² KnowBe4 Research. *Phishing Threat Trends Report 2025*. KnowBe4, 2025.

³ VIPRE Security Group. *Email Threat Trends Report Q2 2024*. VIPRE Security Group, 2024.

⁴ MIT Sloan School of Management and Harvard University. *Research on AI-Generated Phishing Effectiveness*. 2024.

Variability as a defensive challenge

AI did not just improve language quality. It eliminated predictable repetition.

Historically, phishing campaigns reused templates at scale. Detection systems could cluster similar messages or identify signature reuse.

In 2025, a single sender domain could distribute invoice lures, voicemail alerts, fax notifications and secure document requests — each structurally distinct but operationally linked. The infrastructure remained constant while the messaging continuously changed.

This variability reduced the effectiveness of template-based detection and signature clustering.

From content signals to intent signals

As surface-level indicators declined, the malicious signal shifted inward.

Modern phishing messages often:

- Authenticate properly
- Use legitimate cloud infrastructure
- Contain no obvious malware
- Avoid grammatical errors

The risk lies not in how the message looks, but in what it attempts to induce: payment diversion, credential reauthentication, urgent verification or cross-channel engagement.

This marks a structural transition in email security — from identifying malformed content to modeling behavioral intent.



AI on both sides of the inbox

As offensive AI capabilities matured, defensive AI evolved in parallel.

In 2025, INKY expanded its use of AI and GenAI to evaluate email through contextual analysis rather than isolated technical indicators. Detection increasingly incorporates signals such as:

- Behavioral intent modeling
- Sender communication pattern analysis
- Computer vision–based rendering consistency
- Multi-signal classification across message attributes

GenAI assigns contextual labels where appropriate, reflecting the layered structure of modern phishing campaigns.

The goal is not simply to identify obvious anomalies. It is to determine whether a message aligns with expected behavioral patterns. Where attackers use AI to generate variation, detection must model deviation. Where surface indicators disappear, analysis must shift toward intent.

Real-time user guidance inside the inbox

Even with advanced detection, user decisions remain a critical part of the security equation.

Phishing campaigns increasingly rely on behavioral manipulation rather than technical compromise. As a result, security controls must support users at the moment decisions are made.

Throughout 2025, INKY continued to refine its in-inbox guidance model, improving banner clarity, contextual explanations and one-click reporting capabilities. Rather than presenting generic warnings, banners provide immediate context about why a message may be risky and encourage verification through safer channels.

This approach reframes the inbox from a passive delivery channel into an active layer of risk mitigation, helping interrupt phishing attempts before they escalate into account compromise or financial loss.

These developments reflect a broader shift in the structure of email security. As attackers adopt AI, exploit trusted infrastructure, and design campaigns around human behavior, defensive strategies must evolve accordingly. Detection models, user guidance, and operational workflows are all adapting to a threat landscape defined less by obvious technical compromise and more by contextual manipulation. The implications extend beyond individual phishing campaigns. They point to deeper structural changes in how cybercrime will operate in the years ahead.

Predictions for 2026: Escalation, consolidation and AI at scale

Last year, we predicted that generative AI would become essential to detecting future phishing attacks. That prediction proved correct.

In 2025, AI did not merely enhance phishing — it normalized it. Polished, grammatically flawless, highly targeted attacks became the baseline rather than the exception. Detection strategies built around language errors and template repetition rapidly lost relevance.

Looking ahead to 2026, the trajectory is clear: AI will continue to accelerate both offensive and defensive capabilities. But the broader story extends beyond email content.

The next phase of cybersecurity evolution will be defined by five forces.



1 AI arms race: From automation to human-level analysis

Generative AI has permanently lowered the barrier to producing convincing phishing campaigns. What once required expertise and time can now be generated instantly, enabling highly personalized, context-aware attacks at scale.

In 2026, phishing will become even more targeted, variable, and linguistically polished — further reducing the usefulness of surface-level detection signals.

At the same time, organizations are accelerating adoption of AI-driven defense.

Security models are shifting from static rules toward reasoning-based systems that evaluate intent, behavioral consistency and contextual anomalies.

The next phase of email security will not be defined by filtering alone, but by AI systems capable of analyzing messages holistically and adapting continuously as tactics evolve.

The arms race will intensify, but defensive AI will mature alongside offensive AI.

2 Infrastructure risk will eclipse endpoint risk

While phishing remains the dominant entry point, attackers are increasingly targeting shared platforms rather than individual organizations.

SaaS providers, identity platforms, cloud infrastructure and widely deployed security tools present attractive targets because compromise at that level scales instantly across ecosystems.

The concentration of digital infrastructure among a small number of hyperscale providers introduces systemic risk. A breach within a widely adopted platform could have cascading effects far beyond a single enterprise.

In 2026, the industry may confront incidents that test not just individual security programs, but the resilience of interconnected ecosystems.

3 The talent gap meets AI acceleration

As threats grow more complex, the cybersecurity workforce is undergoing structural change.

AI is automating many entry-level technical functions, potentially compressing the traditional pathway through which analysts gain experience. This raises an important question: if large-scale incidents occur, will organizations have enough experienced responders?

AI will augment detection and triage, but human judgment remains essential in high-stakes incident response and strategic defense.

In 2026, cybersecurity risk may be shaped as much by workforce dynamics as by technical vulnerability.

4 MSPs enter a strategic transition

For Managed Service Providers, 2026 may mark a transition year.

AI and automation are becoming foundational to both operational efficiency and competitive differentiation. MSPs that embed AI into workflow automation, threat detection and advisory services will improve margins while delivering stronger security outcomes.

Those that delay adoption risk increased operational cost and reduced strategic relevance.

The opportunity is not limited to internal efficiency. SMB customers will increasingly look to MSPs for guidance on AI-driven transformation. Providers that lead by example will define the next phase of the market.

5 The evolution of email security

Email will remain central to cybercrime, but its role will continue to evolve.

In 2026, organizations should expect:

- Expanded abuse of collaboration and SaaS platforms
- Growth in calendar-based and cross-channel phishing
- Increased impersonation realism driven by generative AI
- Continued shift toward intent-based detection models

Phishing will increasingly resemble legitimate communication. Detection must increasingly rely on contextual reasoning rather than isolated indicators.

The direction is clear: email security will become less about identifying known threats and more about continuously modeling behavior and intent.



Closing outlook

The trends outlined in this report reflect a structural shift in how email-based attacks are evolving.

Phishing has evolved from opportunistic spam into a precision instrument of fraud. AI has lowered the barrier to producing highly convincing attacks, while the concentration of infrastructure has increased the potential blast radius of compromise. At the same time, organizations are accelerating adoption of AI-powered defense to keep pace.

The result is not a simple escalation narrative. It is an inflection point.

In 2026, cybersecurity will be shaped less by isolated techniques and more by systemic forces: automation on both sides, infrastructure interdependence and the pace at which defenders can adapt detection models to changing attacker behavior.

Email will remain central to this dynamic — not because it is outdated, but because it is universal. Its ubiquity makes it both indispensable and persistently exploitable.

Understanding that reality and building security strategies that assume continued evolution rather than stability will define resilience in the year ahead.

Stop email threats before they reach your users

Kaseya's GenAI-powered email security solution, INKY, continuously scans emails to detect and prevent phishing attempts. It neutralizes harmful messages before they can compromise user accounts or infiltrate your organization.

[Learn more](#)



Kaseya is the leading global provider of AI-powered IT management and cybersecurity software. Kaseya delivers a unified technology platform to manage infrastructure, secure endpoints, back up critical data, and streamline operations for more than 40,000 MSP and SMB customers around the globe. To learn more, visit www.kaseya.com.

[kaseya.com](https://www.kaseya.com)