Inbox intelligence

White Paper

Five email threats bypassing company email security

Table of contents

Executive summary

~50%

Share of phishing emails in late 2025 containing AI-generated content

Hoxhunt Phishing Trends Report 2026

27%

Share of phishing detections that pass SPF, DKIM, and DMARC authentication

Group-IB Business Email Protection — H1 2026 detections

80%

Share of phishing detections with no malicious attachment — caught by content, sender, and URL analysis alone

Group-IB Business Email Protection — H1 2026 detections

Email remains the dominant initial-access vector for cyberattacks. Despite decades of investment in authentication, gateways, sandboxes, and user training, attackers continue to find email the most reliable channel to reach their targets. What has changed over the past few years is the scale and sophistication. Phishing now operates as an industry, with commercial kits sold by subscription. Multi-factor authentication can now be bypassed automatically. On top of that, AI tooling has lowered the cost of an effective attack and raised the ceiling of how convincing one can be.

This report covers five email attack patterns that defined late 2025 and early months of 2026. Each pattern bypasses the security controls most organisations rely on. Two are technique-driven: Adversary-in-the-Middle phishing kits and adaptive Phishing-as-a-Service kits, which have reshaped how phishing operates at scale. One is delivery-driven: callback phishing, which removes everything conventional email security relies on for detection. One targets a specific class of victim: developers of open-source maintainers in supply-chain attacks, where email authentication's own design is used to bypass it. Last but not least: Payload-driven infostealer email campaigns, which fuel the ransomware and business email compromise (BEC) economy.

Author

Written by Group-IB specialist

Anton Shumakov

Machine Learning Engineer

Attack 1

Adversary-in-the-Middle phishing kits

Adversary-in-the-Middle phishing places a reverse proxy between the victim and the legitimate login service. The proxy captures both credentials and the session cookie issued after multi-factor authentication (MFA), making MFA irrelevant. The technique is now widely available through Phishing-as-a-Service kits including Tycoon 2FA, Evilginx, EvilProxy, and Sneaky 2FA.

What makes AiTM effective is that it does not break MFA — it simply waits for the victim to complete it correctly. The fake login page is pixel-perfect because it is a live proxy of the real one. The victim's password, the MFA prompt, and the session cookie all flow through the attacker before reaching the legitimate service. From the victim's perspective, the login succeeds normally. From the security stack's perspective, the authentication completes with valid credentials and a valid MFA approval. The compromise is invisible until the attacker uses the stolen session.

Real-world case — a payment lure that hijacks a Microsoft 365 session

The lure is a payment notification with the subject “Transfer Processed: Your Remittance Receipt”. The body is brief: a transaction has been processed, with a “View Remittance Receipt” button beneath. The sender is typically a compromised legitimate account or a spoofed business identity. There is no urgency in the language, no attachment, and no embedded threat for a content classifier to flag. Financial recipients open payment confirmations as part of their normal workflow, and the email is built to resemble one of them.

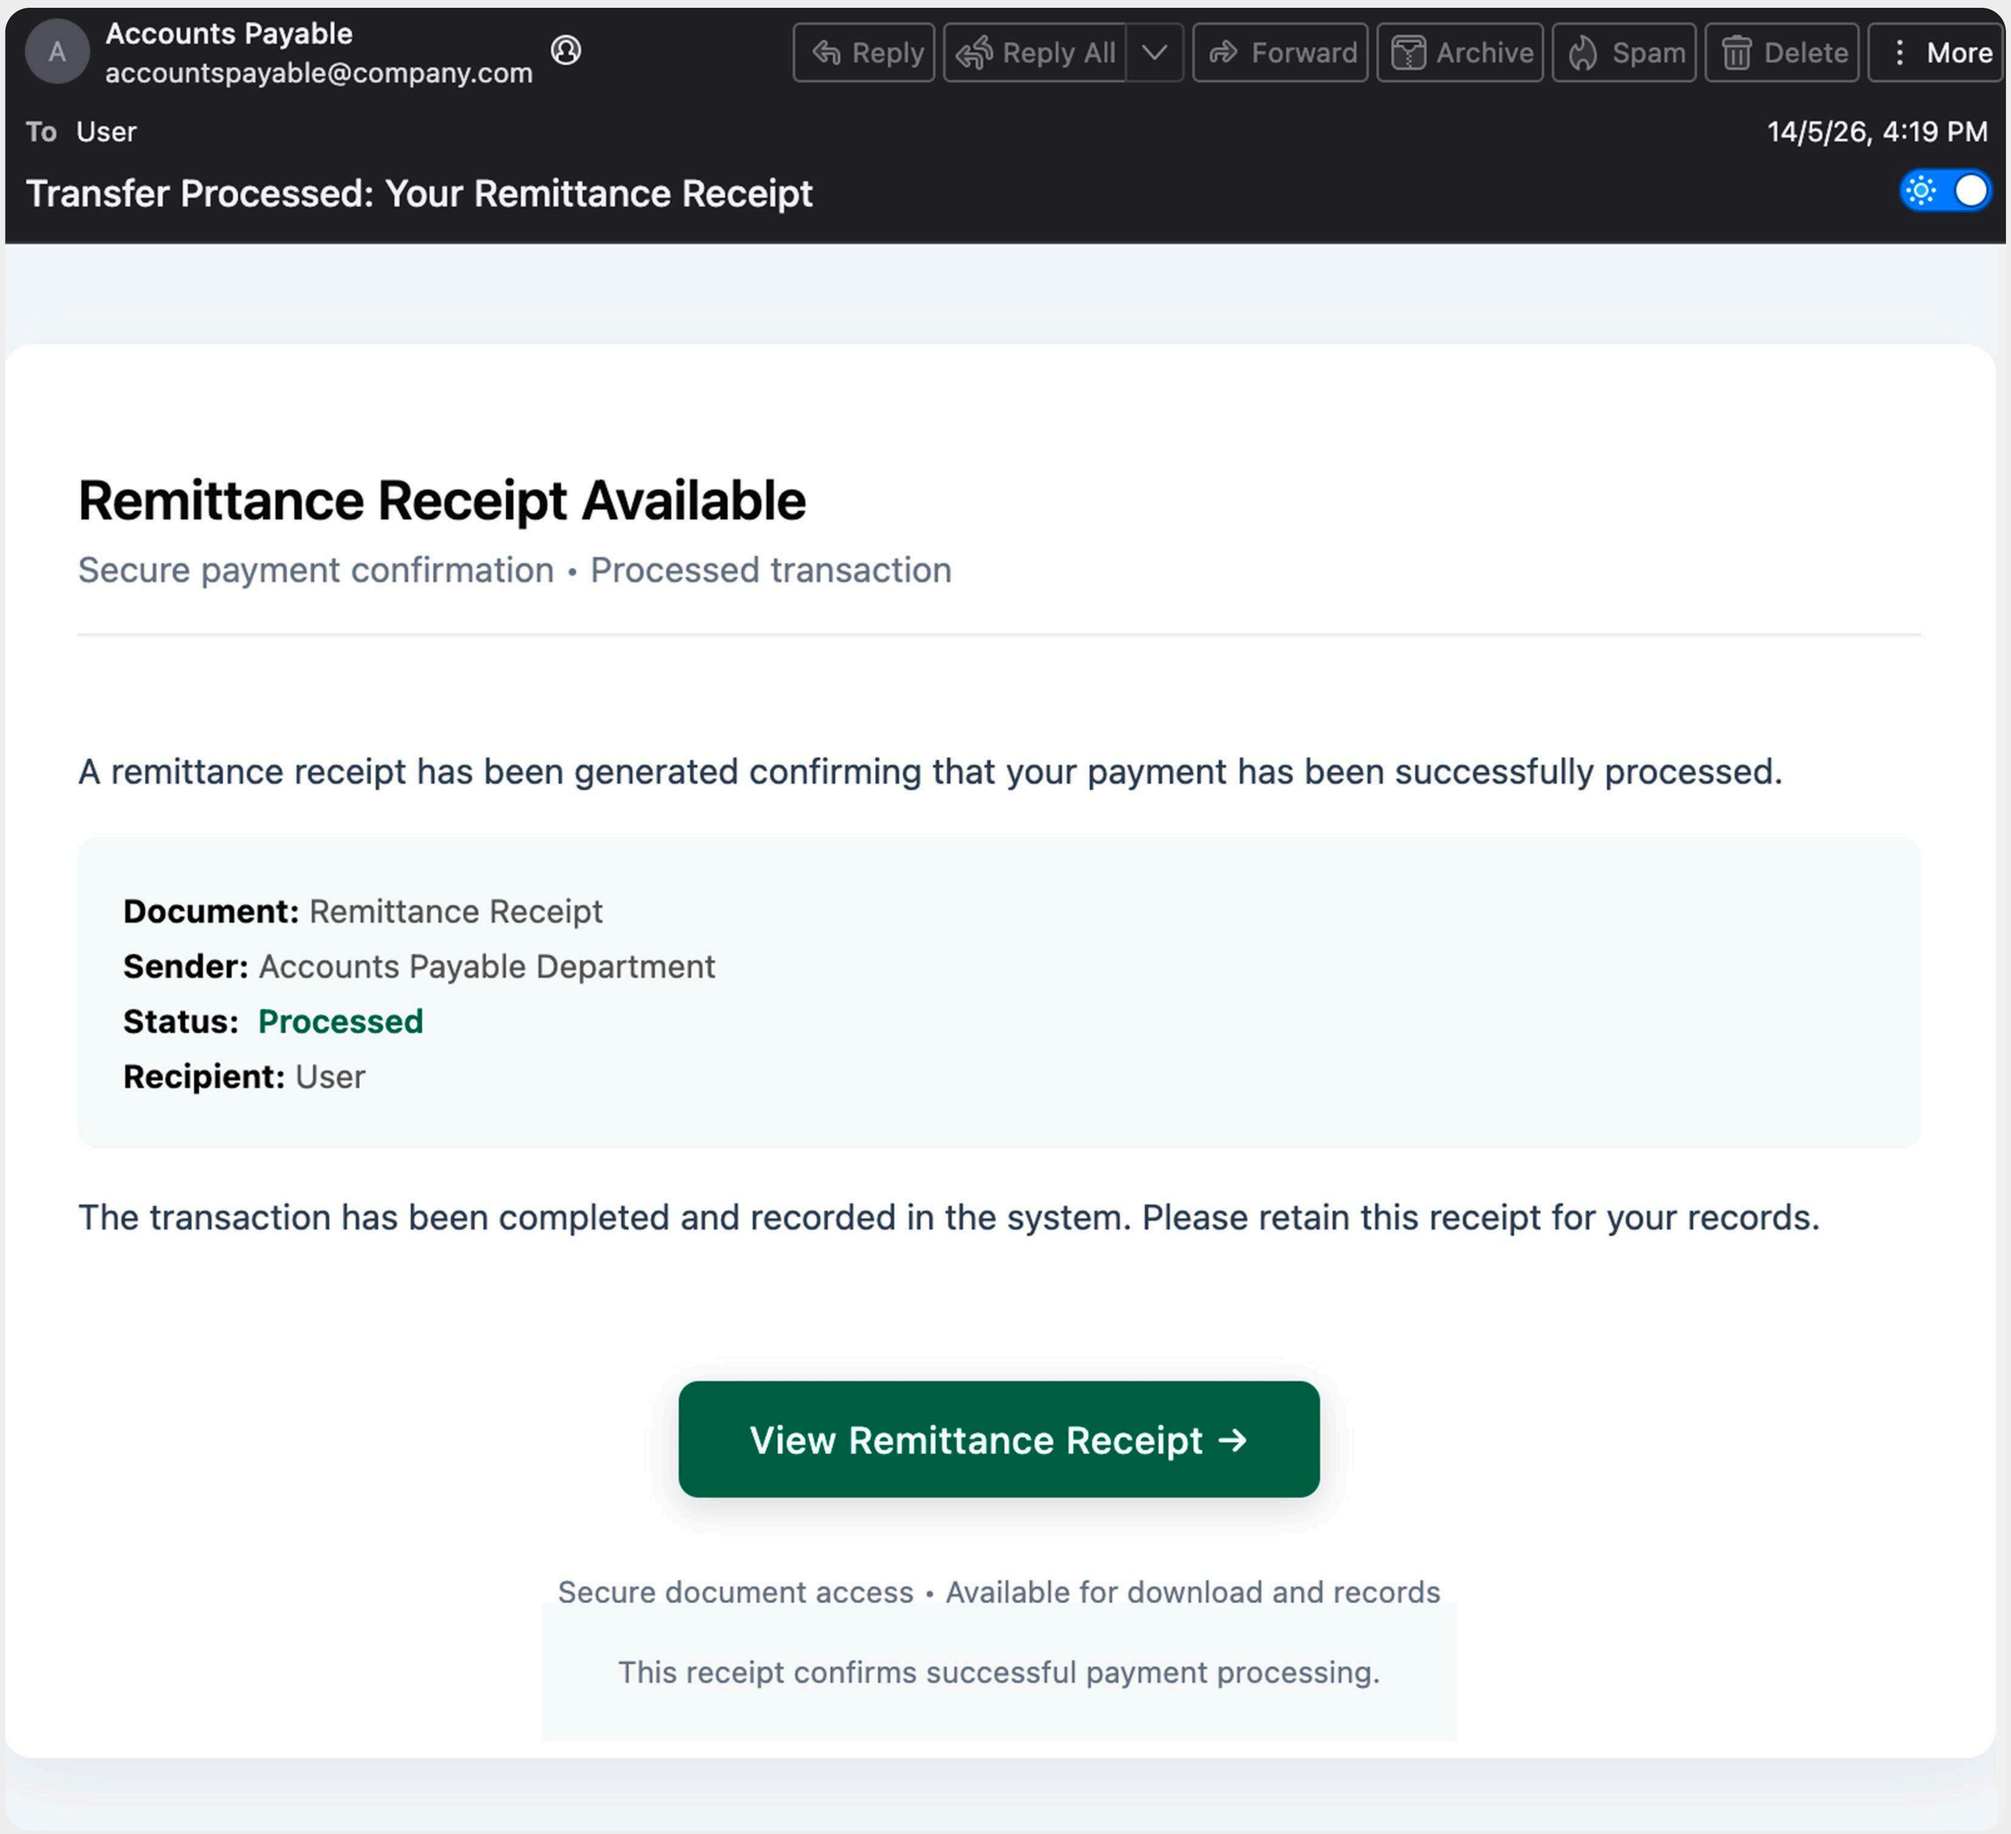

Clicking the button triggers a series of redirects through legitimate SaaS platforms and short-lived domains before reaching the actual AiTM proxy. A CAPTCHA gate, typically a fake "HumanCheck" challenge, filters out automated scanners before displaying the login page. The proxy then presents a Microsoft 365 login page identical to the legitimate one, with the victim's email address pre-filled.

The victim enters their password and completes the MFA prompt. The proxy relays the credentials and authentication request to Microsoft, which authenticates the session and issues a session cookie. The proxy captures the cookie while maintaining the appearance of a normal login process. From the victim's perspective, the login succeeds as expected. The attacker now possesses a valid authenticated session that can be used without repeating the MFA challenge.

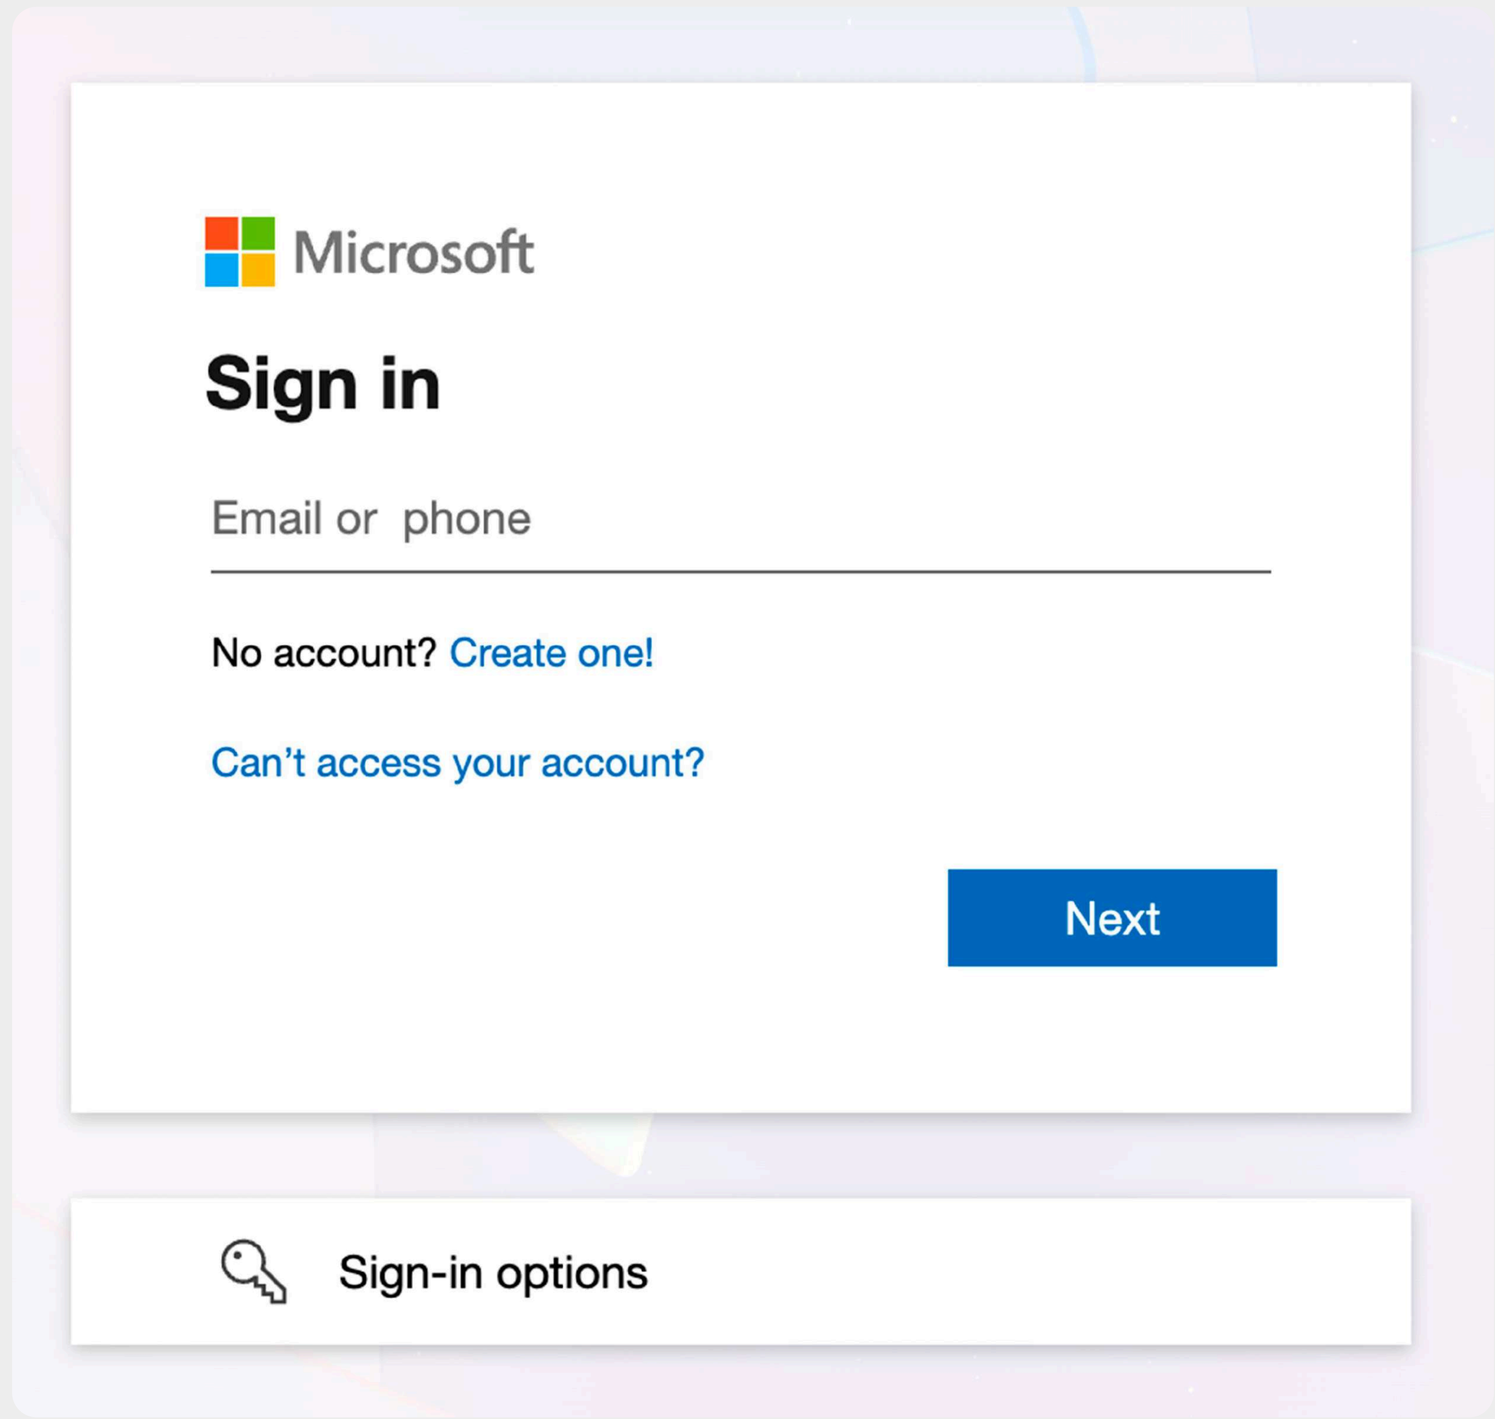

Detection: correlating the redirect chain

The sender is spoofed, and the recipient has no prior correspondence history with the address. The intermediary domains in the redirect chain are newly registered, often only hours before the campaign begins, which can be confirmed through RDAP records. The redirect path itself is unusual: a payment notification passing through multiple legitimate SaaS platforms and short-lived domains does not resemble a legitimate remittance workflow. Individually, these indicators may appear benign. Together, they reveal a coordinated phishing chain.

The final page provides the clearest evidence. After passing through a CAPTCHA gate, the victim reaches a page that renders as a Microsoft 365 login portal, identifiable by its DOM structure, branding elements, and credential-capture form. A Microsoft login page appearing at the end of a redirect chain initiated by a payment-themed lure, hosted on a newly registered domain unrelated to Microsoft, and pre-populated with the recipient's email address is sufficient to establish phishing intent.

Group-IB Business Email Protection correlates these signals across sender history, redirect-chain anomalies, domain age, and page-content analysis to identify and block the message, rather than relying on any single indicator.

Attack 2

Callback scams (telephone-oriented attack delivery)

> GenAI enables threat actors to craft more convincing messages, target victims with greater precision and orchestrate automated multichannel campaigns that blend email, voice calls and text messages.
>
> Gartner, Cybersecurity Threat: Phishing and BEC, 2026

Callback phishing avoids everything a security tool can scan: no URL, no attachment, no executable. The only payload is a phone number. Once the victim calls, the attacker walks them through installing remote access software or transferring money. Ransomware groups use it as their primary entry point. TOAD-as-a-Service platforms now offer multilingual operators and AI voice clones.

The success of callback scams comes from inverting the trust model. Email security has spent decades training users not to click suspicious links or open suspicious attachments. A phone call sits outside that training. When the email itself contains no obvious threat — just a phone number for questions about an invoice — the user's instinct is to call rather than ignore. Once the conversation moves to voice, the attacker has a real-time, high-trust channel where social engineering can be far more aggressive than over email. By moving the interaction from a company-managed device to a personal phone call, the attacker also circumvents enterprise security controls entirely.

Real-world case — a fake Norton renewal receipt

The email arrives as a Norton service receipt confirming a "Premium Device Security Plan" the recipient never purchased. It includes a receipt number, an activation date, a twelve-month subscription period, and a three-figure charge, presented like a legitimate invoice. There is no link or attachment. The only way to act is to call a support number, displayed prominently near the top of the message and repeated in the footer.

For someone who does not recognize a charge they never authorized, the natural reaction is to call. The number connects to an operator posing as Norton support. The caller is guided through a “cancellation” process designed to harvest credentials and one-time passcodes (OTPs), or persuaded to install a remote-support tool under the pretext of reversing the charge. Once remote access is established, the attacker can use the victim's authenticated session to move funds.

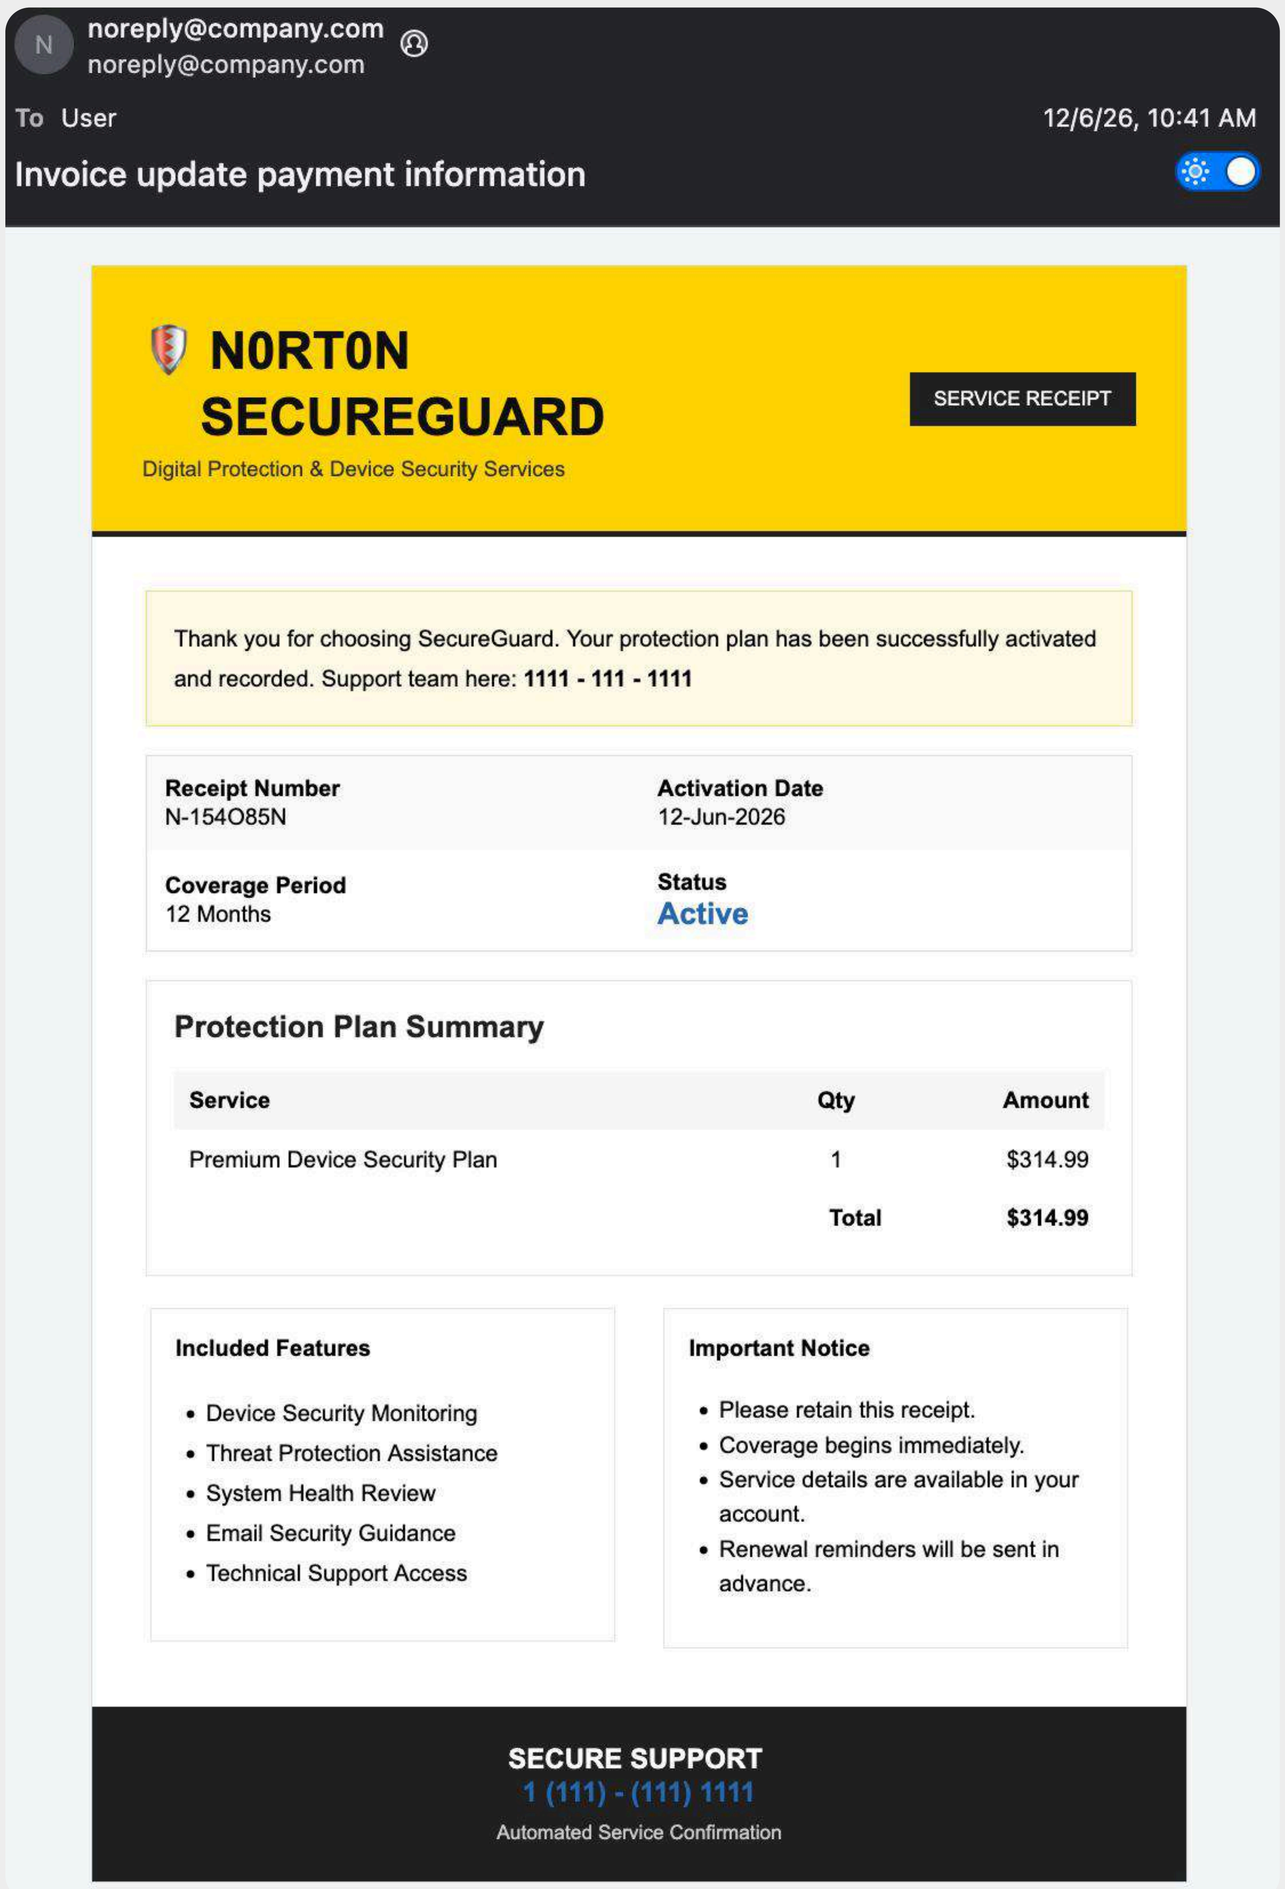

Detection: when clean authentication proves nothing

The signal lies in the combination of indicators, because each one appears legitimate in isolation. Authentication even passes: SPF, DKIM, and DMARC all align. But that is only because the message is sent from a genuine, compromised mailbox rather than a spoofed one, so a clean authentication result proves little.

The indicators emerge one layer deeper. The envelope sender and the From address resolve to different domains, while the Message-ID domain matches neither. The From domain also lacks its own DKIM signature. The message originates from a public, low-trust domain and is relayed through consumer mailing groups that expand a single message into a long recipient list, a delivery path no legitimate vendor receipt would follow.

The content reinforces the technical indicators. The brand name is deliberately misspelled using lookalike characters to evade keyword matching. The message is framed as a payment notification, creates a sense of urgency, and provides a phone number as the only way to dispute the charge, with no link or attachment.

Group-IB Business Email Protection correlates these signals rather than evaluating them individually. It identifies the sender mismatch and Message-ID anomaly as indicators of impersonation, detects the missing DKIM signature, the public-domain origin, the relay path, and the payment-themed social engineering, then correlates those findings with known phishing and bulk-mail activity. The combined evidence identifies the message as malicious.

Attack 3

Software supply chain phishing

Software supply chain attacks target the maintainers of widely used open-source packages. By compromising a single maintainer account, attackers can introduce malicious code into software distributed to millions of users. The attack typically begins with a lookalike domain configured with valid SPF, DKIM, and DMARC records. The maintainer then receives what appears to be a legitimate MFA reset request or security notification. Because the email originates from an attacker-controlled domain that is properly authenticated, all three authentication checks pass, leaving traditional reputation-based controls with little reason to flag it.

The potential impact of a supply chain attack is enormous. A single compromised maintainer can push a malicious update to every system that depends on the affected package. Rather than targeting thousands of downstream organizations individually, attackers need only compromise one person with publishing privileges to a widely used library. For popular packages, the impact can extend to hundreds of millions of installations across financial services, government, healthcare, and developer ecosystems.

2.8B

Combined weekly downloads of the 20 NPM packages compromised in the npmjs.help campaign

Group-IB CERT research, October 2025

Real-world case — the npmjs.help campaign

Attackers registered the domain npmjs.help and configured it with valid SPF, DKIM, and DMARC records. Phishing emails were sent from support@npmjs[.]help, framed as a required two-factor authentication update with the threat of account suspension. The credential-harvesting page mirrored the real NPM login interface precisely. Once credentials were entered, the attackers logged in and pushed a malicious JavaScript clipper into twenty popular npm packages with a combined 2.8 billion weekly downloads. The clipper monitored cryptocurrency wallet activity and silently replaced wallet addresses with attacker-controlled ones.

At least four npm developers received identical emails in a coordinated wave. The emails passed all three authentication checks because the attacker had configured their own domain correctly — the domain was newly registered and designed to look like NPM's own infrastructure, not a spoofed legitimate domain. Standard sender reputation checks found nothing to flag.

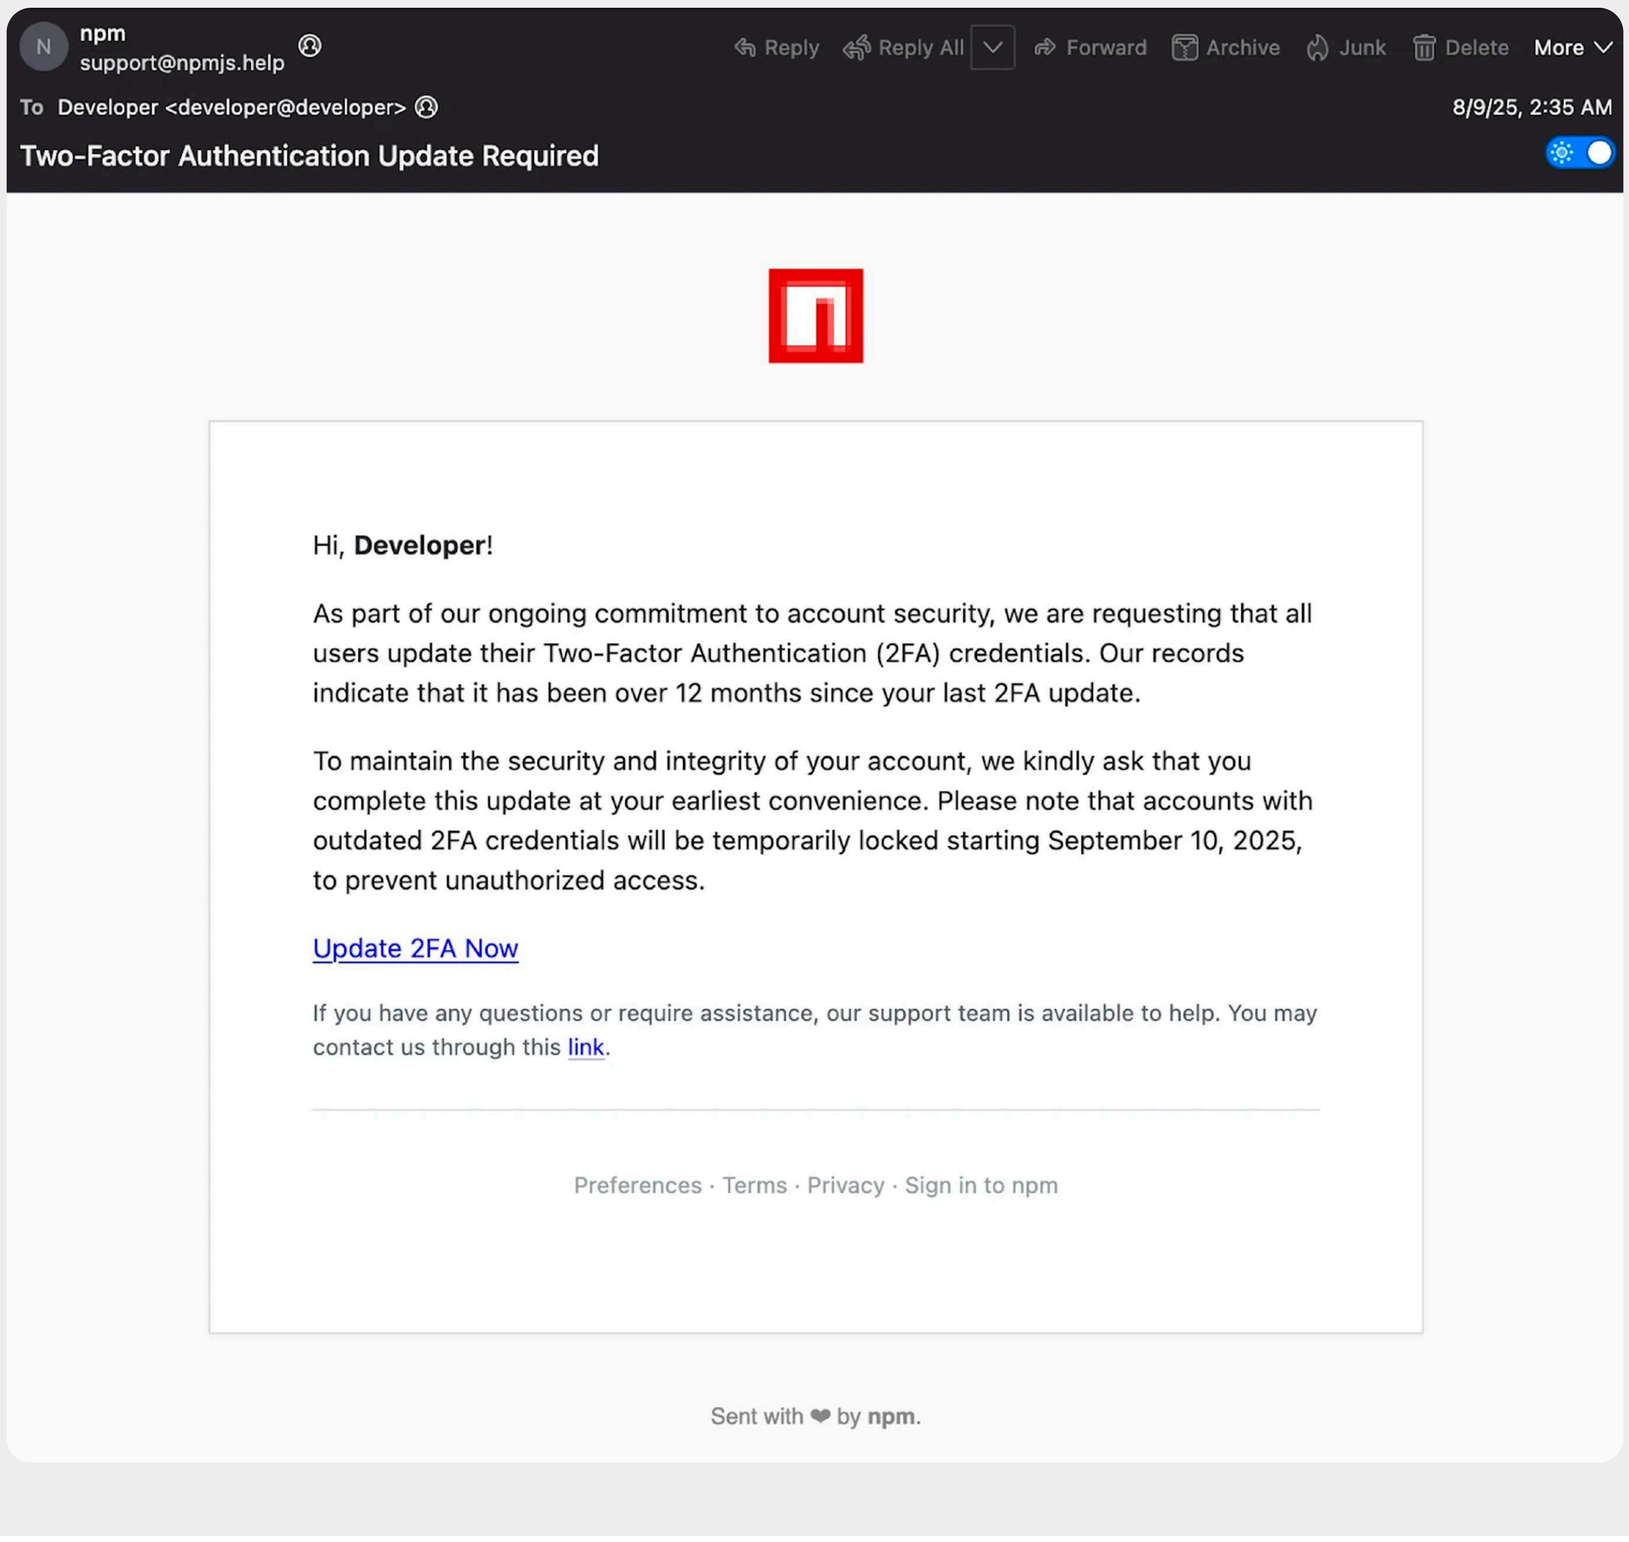

Detection: why authentication is not the answer

The email passes SPF, DKIM, and DMARC because it originates from the attacker's own domain, not a spoofed one. Reputation filters have nothing to flag — the domain is new and unknown rather than known-bad. The email content is security-themed and urgent, which is exactly the language a legitimate MFA notification uses.

Detection requires signals beyond authentication. **Group-IB Business Email Protection** flags the newly registered domain using RDAP analysis, identifies brand impersonation through lookalike-domain matching against NPM's known infrastructure, analyses urgency language and account-suspension framing as social engineering patterns, and detonates the linked credential-capture page in a sandbox to confirm its harvesting behaviour. The domain's age, its structural resemblance to NPM's actual domain, and the content's engineering toward immediate action are together sufficient to block the email before any developer's credentials are entered.

Attack 4

Adaptive phishing kits

Adaptive phishing kits, sold as Phishing-as-a-Service, assemble credential-capture pages in the victim's browser at click time rather than in advance. A kit like LogoKit pulls the target organisation's logo from public services, fetches a background screenshot of the website, and builds a login page that mirrors the real one — customised per victim based on the email address passed in the URL. Because the page is assembled only when clicked, URL scanning at delivery sees a benign static page. Kits are hosted on IPFS or cloud storage to complicate takedowns.

The economic appeal of adaptive kits is that one infrastructure base serves any number of targets. An attacker does not need to build a separate phishing page for each company they want to impersonate. The kit fetches the company's branding live — via Clearbit for favicons and thum.io for website screenshots — and generates the page automatically based on the recipient's email domain. Static URL signatures and page-hash blacklists have limited effect: the URL stays the same across many targets, but the rendered page changes for each one.

22%

Year-over-year growth in phishing activity through 2025

Group-IB High-Tech Crime Trends Report 2026

Real-world case — DocuSign impersonation wave using LogoKit

Since late August 2025, Group-IB has observed a sustained wave of DocuSign impersonation emails using the LogoKit phishing framework. The emails closely follow the legitimate DocuSign notification design. The sender is spoofed — sometimes mimicking the recipient's own domain — and the message addresses the user by their login name. The Reply-To header frequently mismatches the sender, redirecting to an unrelated company or a generic Gmail account.

The “Review document” button leads to a URL hosted on an IPFS gateway or AWS S3 bucket, with the target's email address embedded in the URL. When the link is opened, LogoKit uses that email address to fetch the recipient's organisation branding live — a background screenshot from thum.io, a favicon from Clearbit — before assembling a customised login screen designed to harvest credentials. From the victim's perspective, the page looks exactly like their own organization's login portal. Every victim sees a different page, yet every page is built from the same infrastructure.

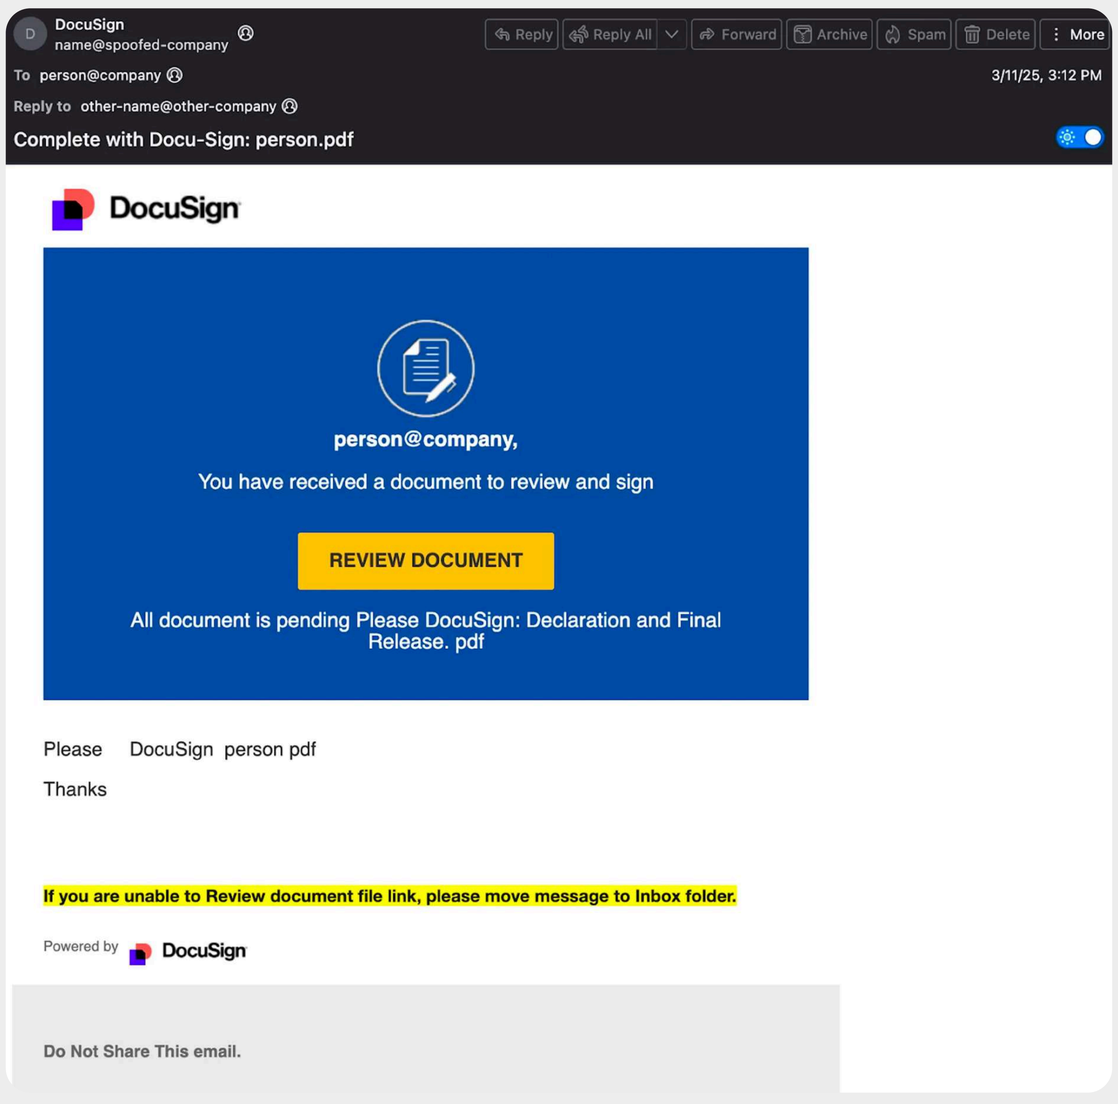

The email shows an SPF failure immediately, due to sender spoofing. Several other indicators are present on inspection: the Reply-To points to an unrelated organisation or a public email provider; the target's email address is embedded as a parameter in the URL; the links direct to an IPFS gateway or AWS S3 bucket rather than DocuSign's own infrastructure. A forged sender, infrastructure unrelated to DocuSign, and hosting patterns characteristic of phishing pages together establish clear intent to harvest credentials.

Detection: from delivery to the final click

Group-IB Business Email Protection detects this campaign by correlating signals across multiple detection layers. Sender and infrastructure analysis identifies the mismatch between the Reply-To and From headers, the failed SPF check, and the absence of a DKIM signature. Content and linguistic analysis detects spam-like patterns in the email body, as well as the recipient's login name in the subject line, a personalization tactic commonly used to increase click-through rates. URL and behavioral analysis flags the inclusion of the recipient's email address in the URL, the use of IPFS or AWS S3 hosting, and prior associations between the URL and known spam campaigns.

In cases where the URL appears benign at the time of delivery because the attacker has deliberately "disarmed" the link during the delivery window, Group-IB's Time-of-Click analysis evaluates the page in its live, rendered state. When the user clicks the link, the system re-evaluates the destination and extracts relevant indicators, including hosting details, URL parameters, and dynamically loaded content. If the page generates a credential-capture form or exhibits brand impersonation behavior, Business Email Protection immediately blocks access and displays a warning page.

This approach allows the attack to be stopped at the point of interaction, even when earlier indicators are minimal or absent.

Attack 5

Infostealer email campaigns

Infostealers are commercial credential-theft toolkits (e.g., Phantom Stealer) sold as subscriptions and distributed through email. The lures are mundane business correspondence: procurement requests, invoices, shipping notifications. Attachments are archives (.tar, .7z, .zip, .rar) containing a dropper that installs the stealer. From there, the stealer collects passwords, session cookies, cryptocurrency wallets, and authentication tokens, exfiltrating through Telegram or Discord.

The infostealer economy is largely commoditised. Stealer logs (i.e., bundles of credentials, cookies, autofill data, and crypto wallets harvested from a single victim) are traded openly on underground marketplaces. Initial access brokers buy logs in bulk, sort them by victim profile, and resell access to ransomware operators or fraud crews. A single employee's stealer log can include corporate VPN credentials, SaaS session cookies, and personal banking data in the same file. This is why infostealer delivery via email is not a credential theft problem in isolation — it is the first stage of ransomware, BEC, and supply chain compromise. The constant resale value of these logs is what keeps the campaigns running. The lures are deliberately mundane because anything more elaborate raises suspicion, and mundane content is harder for classifiers to flag.

2,227

Publicly advertised corporate access sales on dark web forums detected by Group-IB in 2025. The most valuable types — domain admin, VPN panels, SaaS logins — now trade exclusively in private channels and never reach open markets

Group-IB High-Tech Crime Trends Report 2026

Real-world case — Phantom Stealer five-wave campaign

Between November 2025 and January 2026, Group-IB observed five waves of Phantom Stealer delivery targeting European logistics, manufacturing, and technology organisations. Every email originated from a spoofed sender impersonating a legitimate equipment trading company, with procurement-themed subject lines designed to resemble routine supplier correspondence. Each email carried an archive attachment containing an obfuscated JavaScript dropper.

The five waves targeted multiple unrelated organisations on the same days — a pattern consistent with stealer-as-a-service operations running coordinated campaigns rather than targeted attacks. Across all five waves, every email was blocked by Group-IB Business Email Protection before reaching end users.

The malware itself collected browser credentials, cookies, saved passwords, autofill data, and payment card information, and extracted session data from messaging and email platforms, Wi-Fi credentials, and authentication tokens — exfiltrating through Telegram. Once in circulation on underground markets, this data fuels downstream ransomware initial access, BEC fraud, and supply chain compromise. A single employee credential from one of these campaigns can become the entry point for an incident measured in millions of euros.

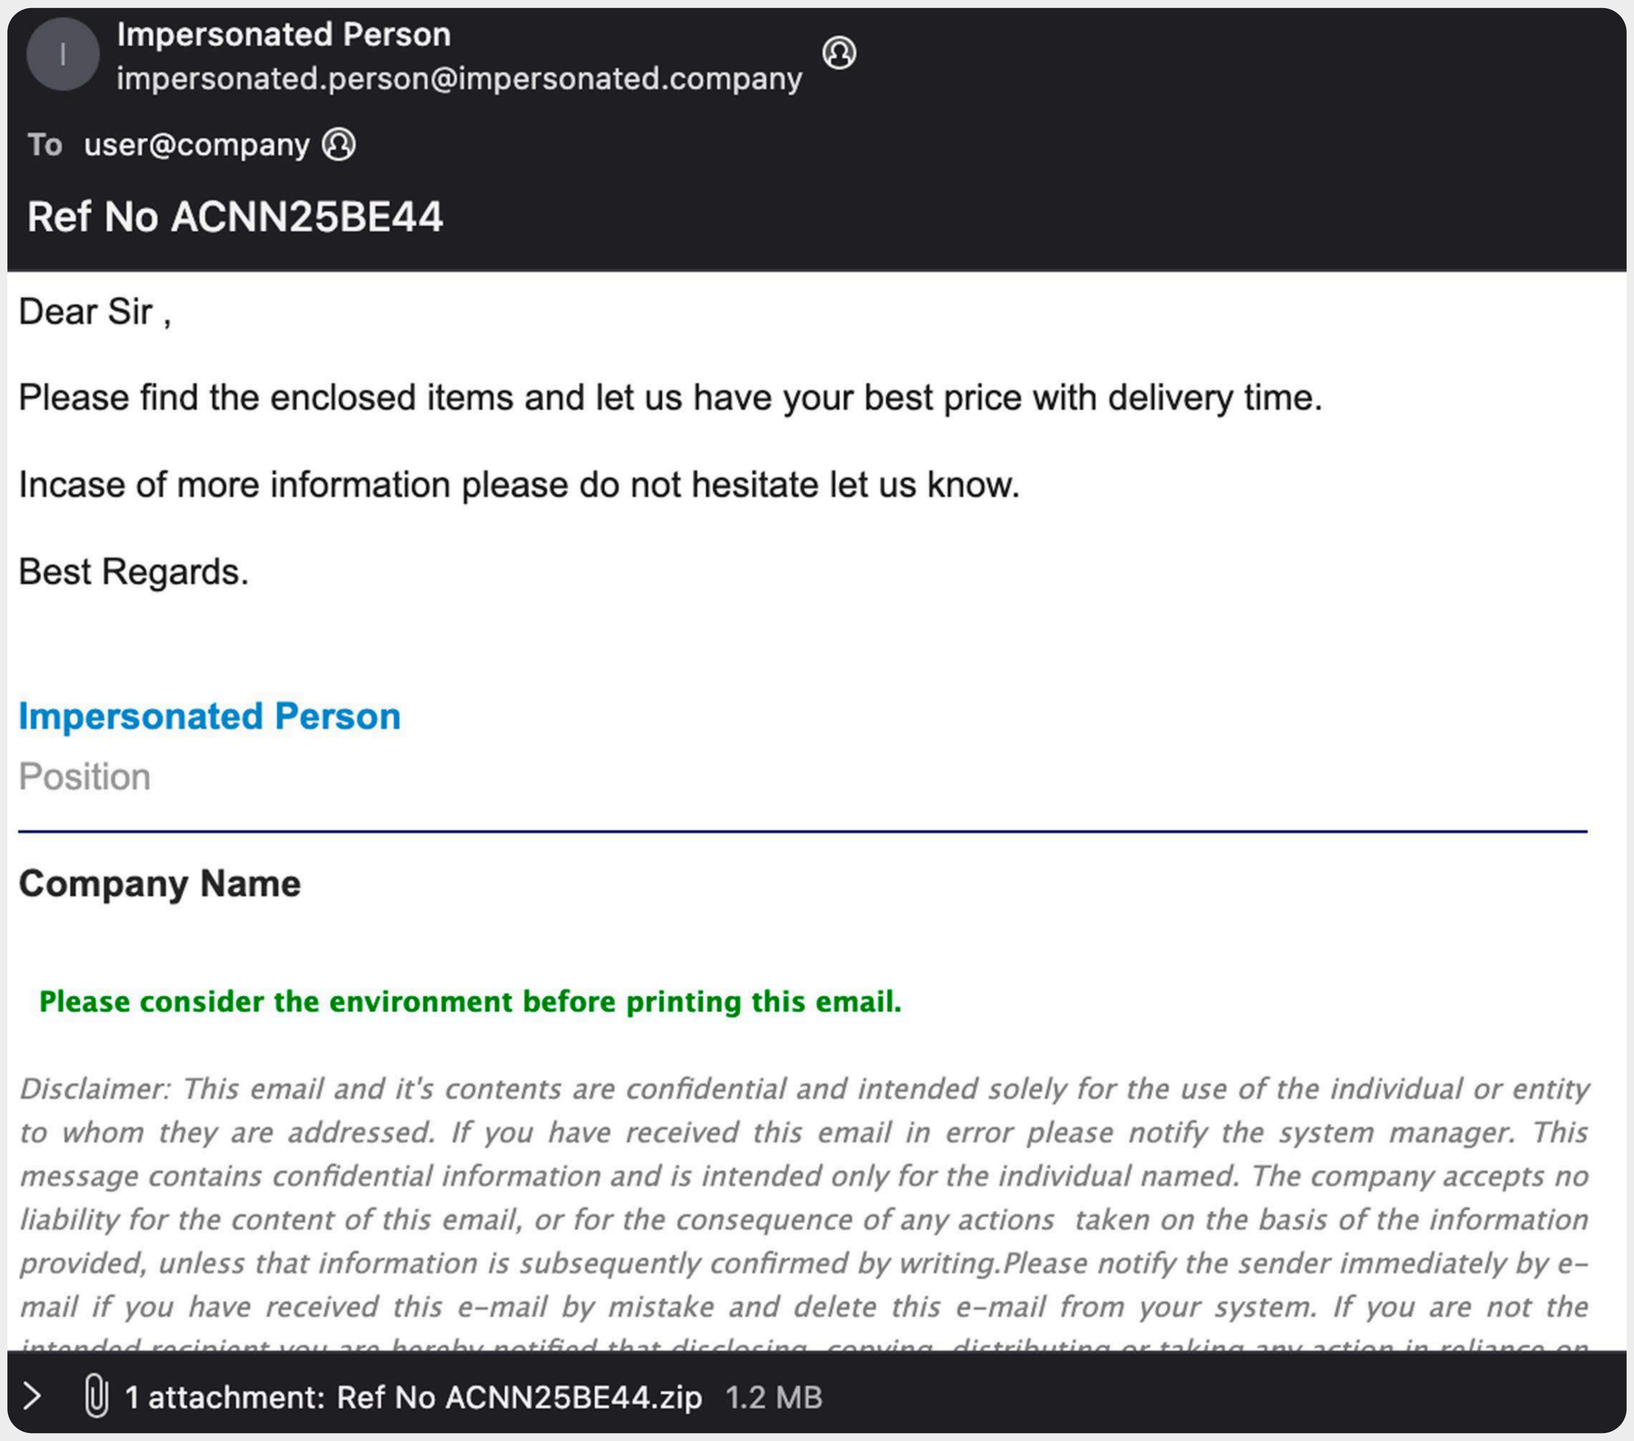

Detection: before and after detonation

Detection can happen before detonation is needed. The emails show SPF failures, missing DKIM signatures, and Message-ID alignment patterns characteristic of impersonation tooling. The procurement-themed subject lines are reused across waves with consistent linguistic fingerprints. Sender impersonation is detectable through domain analysis even before the attachment is opened.

For cases where runtime analysis is required, **Group-IB's Malware Detonation Platform** — part of Business Email Protection — mimics the customer's actual endpoint environment rather than a generic sandbox. This matters because Phantom Stealer, like most commercial infostealers, is sandbox-aware: it detects generic detonation environments and suppresses its payload. A customer-environment-specific detonation triggers the full execution chain, confirming credential harvesting behaviour and anti-analysis evasion techniques before the email reaches any inbox.

Where email threats are heading

Agentic phishing will make attacks self-optimizing

AI agents are beginning to change the economics of phishing campaigns. Rather than a human operator sending a fixed lure, agentic systems will test multiple subject lines, sender configurations, and attachment formats against a target before selecting the highest-performing variant. The first email the victim sees will already be the version that passed internal quality checks. Combined with real-time data scraping from LinkedIn, company websites, and previous breach data, these agents will produce lures that are contextually accurate at an individual level — not just personalised with a name, but referencing actual business relationships, recent transactions, and internal terminology. The detection challenge this creates is that the attack surface shifts from identifying known patterns to identifying novel combinations that have never appeared before.

Credential theft will tighten its loop with BEC

The pipeline from credential theft to business email compromise is accelerating. As credential marketplaces automate the matching of stealer logs to target organisations and initial access brokers compete on speed, the window between compromise and exploitation is compressing. Organisations that rely on detection inside the mailbox — behavioral anomalies, content analysis, identity signals — will face attackers who have already operated undetected long enough to understand communication patterns, identify financial workflows, and select targets for BEC. The most effective intervention point is before the first login, which requires visibility into the underground markets where credentials are traded before they are used.

The gap between authentication and actual security will widen

The five attacks in this report share a structural characteristic: all of them exploit the trust that email security infrastructure places in correctly authenticated messages. AiTM kits bypass MFA by capturing session cookies after correct authentication. Supply chain phishing uses properly configured SPF, DKIM, and DMARC on attacker-controlled domains. Adaptive phishing kits deliver clean URLs that arm only at click time. Infostealer campaigns use archive formats that sandbox-aware malware detects and evades. As attackers continue to design around authentication rather than through it, the assumption that a correctly authenticated message is a safe message will become increasingly untenable.

How Group-IB Business Email Protection stops these attacks

Group-IB Business Email Protection connects to Microsoft 365 or Google Workspace via API and provides multi-layer protection against the full spectrum of email-delivered threats: phishing, business email compromise, infostealer delivery, adaptive phishing kits, and account takeover.

Unlike solutions that begin detection at the inbox, Business Email Protection is continuously fed by Group-IB's threat intelligence infrastructure to identify compromised credentials in underground markets before attackers use them, and updating detection models based on live campaign tracking.

Group-IB Business Email Protection operates at three layers that most email security platforms cannot reach:

Before the email is sent — threat intelligence

Group-IB Threat Intelligence continuously monitors dark web forums, infostealer log marketplaces, breach databases, and private credential trading channels, with intelligence updated hourly. Business Email Protection checks the organisation's accounts against that intelligence, matching employee email addresses and leaked passwords from stealer logs and breach dumps. When a match is found, the security team receives an immediate alert. The affected password can be reset before any attacker logs in. The BEC that would have followed — the lateral phishing, the invoice fraud, the supply chain compromise — never materialises.

This is a fundamentally different approach from how other email security platforms work. Most derive their compromise signals from what happens inside the email environment: behavioral anomalies, sending anomalies, identity-graph changes. By the time those signals appear, the attacker is already operating inside the mailbox. Group-IB's signal comes from outside the email environment entirely, from the underground economy where credentials are traded, before any email-side indicator exists.

At delivery — multi-signal detection	At delivery, Group-IB Business Email Protection correlates signals across sender authentication, content and linguistic analysis, URL and infrastructure behaviour, and attachment characteristics simultaneously. The 27% of phishing emails that pass SPF, DKIM, and DMARC authentication — and the 80% that carry no malicious attachment — are caught not through any single signal but through the combination of signals each email carries. A spoofed sender with a Reply-To mismatch, a newly registered domain in the URL chain, procurement-themed urgency language, and a recipient's login name in the subject line: any one of these alone might be insufficient to block. Together, they identify the attack at the point of delivery before any user interaction occurs. Files are routed at delivery to an interactive Malware Detonation Platform, which is part of the Business Email Protection solution. Attachments are opened in instrumented environments — macros enabled, nested archives unpacked, password-protected files unlocked using the password recovered from the email body. Behaviour observed during detonation is mapped to MITRE ATT&CK techniques and known malware families, and verdicts feed back into the detection model so future variants are caught by signal correlation rather than detonation.
At click — Time-of-Click analysis	For adaptive phishing kits that deliver clean URLs at delivery and assemble the attack page only when clicked, detection must happen at click time. Group-IB's Time-of-Click analysis evaluates the URL in its live, rendered state at the moment a user clicks — extracting hosting details, URL parameters, dynamic content loading behaviour, and credential-capture form presence. If the page exhibits phishing behaviour at the moment of interaction, Group-IB Business Email Protection blocks access immediately and displays a warning page, regardless of how the URL appeared at delivery.
Internal email security	Once an attacker accesses a mailbox — whether through stolen credentials or a compromised session — every defence designed for external threats becomes irrelevant. Internal emails never cross the perimeter. Most email security tools simply never see them. The attacker has already passed MFA. Every email they send originates from a legitimate account with years of established reputation. Security leaders we surveyed describe this challenge in consistent terms.

Field CISO,
10,000+ employee
SaaS company

After an internal user has been properly authenticated with their MFA, email from them would be trusted. The email would not be subject to the legitimacy inspections like SPF and DMARC since it originates from a bona fide account.

CISO,
mid-sized software
company

This is a traditionally overlooked area due to the assumption that internal communications are intrinsically secure. Don't count on users to sniff out that level of sophistication.

Deputy CISO,
large healthcare
organisation

The phishing messages usually come from a trusted party, which prevents the typical red flags users usually use to identify malicious emails. They can also be sent to hundreds of inboxes at once, which makes it very difficult to quarantine and stop.

Group-IB Business Email Protection connects via API to analyze messages that stay entirely within the organization. If a compromised account sends a lateral phishing email to a colleague, every attachment and link is detonated in a behavioural sandbox and stopped before it reaches the recipient's inbox. Internal email is no longer a blind spot.

See Business Email Protection in action

GROUP-IB

Group-IB is a creator of cybersecurity technologies to investigate, prevent and fight digital crime.

1,600+
Successful investigations of high-tech crime cases

500+
Employees

60
Countries

$1 bln+
Saved by our client companies through our technologies

#1*
Incident Response Retainer vendor

*According to Cybersecurity Excellence Awards

11
Unique Digital Crime Resistance Centers

Global partnerships

- INTERPOL
- EUROPOL
- AFRIPOL

Recognized by top industry experts

- Forrester®
- datos INSIGHTS
- kuppingercole ANALYSTS
- Gartner®
- IDC
- FROST & SULLIVAN

Fight against cybercrime

GROUP-IB.COM
INFO@GROUP-IB.COM

APAC
+65 3159 4398

EU & NA
+31 20 226 90 90

MEA
+971 4568 1785