

**Kaseya®**

WHITEPAPER

# Understanding phishing: Account takeover 2026 update



# The hidden engine behind modern email fraud

One of the most damaging forms of phishing today involves account takeover (ATO) attacks, in which a threat actor compromises a legitimate email account and then uses that access to impersonate a trusted person or organization. These attacks routinely lead to fraud, data exposure and downstream financial loss. Yet they are among the most difficult for traditional secure email gateways (SEGs) to detect.

The FBI categorizes many of these incidents under business email compromise (BEC), which represents the financial outcome of account takeover activity. According to the FBI's 2024 Internet Crime Report, BEC-related attacks accounted for \$2.8 billion in reported losses, with an average loss of \$129,193 per incident. Independent data also shows that wire transfer BEC attacks climbed **over 30% year-over-year in early 2025** and that organizations view BEC as a significant share of their cyber risk landscape. While BEC describes the business impact, account takeover is often the enabling technique — the moment attackers gain control of a legitimate email identity and weaponize trust.

In this report, we examine how attackers obtain and exploit compromised email accounts, why account takeover phishing routinely bypasses legacy email defenses, and how INKY detects these attacks using advanced sender profiling, stylometry and machine learning techniques. Through real-world examples, we show how compromised accounts become high-trust delivery mechanisms for some of the most destructive phishing campaigns organizations face today.

Independent data also shows that wire transfer BEC attacks climbed over

# 30%

year-over-year in early 2025



# High-value email accounts: Why attackers target them

What is an email account worth to an attacker? At first glance, perhaps not much. With more than 8 billion email accounts in use globally, the sheer volume might suggest that any single account has limited value. However, this figure masks an important distinction: individuals and employees routinely maintain multiple email accounts, and not all accounts carry the same level of trust or access.

A randomly created consumer email account may have limited utility. In contrast, a high-reputation corporate, government or academic account represents a powerful phishing asset. These accounts inherit the trust of their domain, their organization, and often the individual behind them. When compromised, they allow attackers to bypass reputation-based defenses and reach targets with messages that appear legitimate by every traditional signal.

In this report, we focus on what happens when credentials are stolen — how attackers leverage compromised accounts to conduct highly effective phishing campaigns that are both difficult to detect and highly damaging.



Attackers use **compromised, high-trust email accounts** to **execute phishing attacks** that routinely evade secure email gateways and exploit human trust.

---

## Account takeover in practice

Every phishing email must originate from a sending account and mail server. Attackers generally obtain this capability in one of three ways:

- Operating their own mail infrastructure
- Abusing shared consumer email services
- Compromising a legitimate account on an organization's mail server

While INKY observes all three approaches, account takeover-based phishing is consistently the most dangerous. When attackers send mail from a compromised corporate account, they benefit from the sender's established reputation, historical legitimacy, and trusted domain. Reputation-based filtering systems are explicitly designed to avoid blocking these senders, as doing so would disrupt legitimate business communication.

Account takeover attacks become even more damaging when attackers impersonate the actual individual who owns the compromised account. Messages appear to come from a known contact, business partner or vendor, making them exceptionally difficult for recipients to question and for automated systems to flag.

## First-party vs. third-party account takeover

Organizations can — and must — protect their own email accounts through strong authentication controls. Multifactor authentication (MFA) is the most effective defense against first-party account takeover and should be universally enforced. Relying on phishing detection alone to protect internal accounts is insufficient when MFA can prevent compromise entirely.

However, organizations have no control over the authentication practices of third parties—vendors, customers, partners, or suppliers whose compromised accounts may be used to target employees. These third-party ATO attacks represent the primary residual risk and require detection after compromise, as prevention is not possible at the source.

Account takeover **enables attackers to evade SEG reputation** filters and convincingly impersonate trusted individuals using their real email accounts.

# How attackers gain access to email accounts

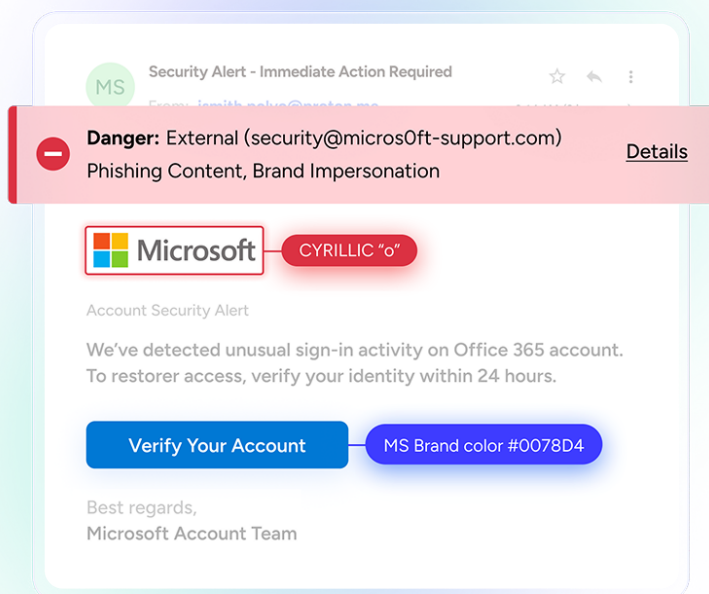
Account takeover begins with a single prerequisite: valid credentials. While there are multiple technical ways to obtain them, attackers consistently favor methods that scale easily and minimize risk. In practice, this means exploiting user behavior rather than attempting to defeat authentication systems directly.

Poor password hygiene remains a contributing factor. Easily guessed or reused passwords can still be compromised through brute-force or dictionary-based attacks, particularly when login rate limiting or lockout controls are weak. Standards bodies such as the National Institute of Standards and Technology (NIST) have long discouraged the use of common or predictable passwords, and modern authentication systems increasingly enforce password strength checks to mitigate this risk.

However, credential harvesting remains by far the most effective and scalable technique. In these attacks, users are tricked into voluntarily providing their credentials through fraudulent login pages that closely mimic legitimate cloud services such as Microsoft 365. These phishing emails often impersonate trusted brands or internal IT notifications and lead victims to attacker-controlled sites designed to capture usernames, passwords, and session data.

Once harvested, credentials can be used immediately or sold, replayed, or combined with other data to enable account takeover. Even strong passwords provide little protection if users are deceived into entering them on a convincing phishing page.

More advanced attackers may also obtain credentials through malware-based techniques such as keylogging or network traffic interception. While effective, these methods typically require prior system compromise and are therefore less common than phishing-based approaches, which require no malware and can be executed entirely through email.



# Conversation hijacking

One of the most effective and difficult-to-detect ATO techniques is conversation hijacking. After compromising an account, attackers often monitor inbox activity — sometimes for days or weeks — waiting for an opportunity to insert themselves into an existing email thread.

This tactic is most frequently observed in high-value transactions, such as real estate closings, invoice payments, or supplier communications. When the timing is right, the attacker replies within an existing conversation, impersonating the legitimate sender and providing fraudulent instructions — often redirecting payments to attacker-controlled accounts.

The effectiveness of this technique lies in its authenticity. The message originates from the real account, appears within an ongoing thread, and references legitimate context. Traditional email security controls struggle to identify these attacks because there is no obvious malicious payload, spoofed domain or anomalous infrastructure.

INKY detects conversation hijacking by identifying subtle inconsistencies between the historical writing patterns of a sender and the content of a hijacked message. Even when headers appear legitimate, stylometric and contextual anomalies can reveal that a message does not match the sender's established profile.



# Detecting third-party account takeover attacks

We've examined how attackers compromise email accounts and how organizations can protect their own users from account takeover through strong authentication controls like multifactor authentication. That leaves a critical gap: third-party accounts.

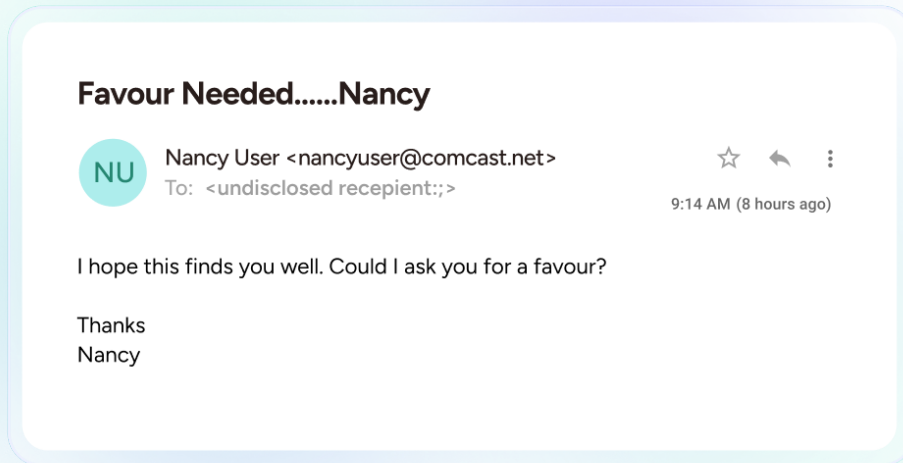
Organizations cannot prevent attackers from compromising the email accounts of vendors, partners, customers or other external contacts. Yet these compromised third-party accounts are routinely used to target employees with phishing emails that appear legitimate by every traditional signal. This makes third-party account takeover one of the most challenging email threats to detect.

This is where a phishing prevention solution like INKY plays a crucial role. While third-party account takeovers cannot be stopped at the source, they can be identified and blocked as malicious emails enter the organization — before they reach an end user's inbox.

To illustrate how INKY detects these attacks, the following sections examine real-world third-party ATO emails that INKY identified and stopped. All examples are shown as received, with only personally identifiable information redacted.



# Why third-party ATO is so difficult to detect



In this example, one of INKY's users received an email from a person named Nancy — a contact she knew and had communicated with previously.

At first glance, there is very little content to analyze. The message is short, lacks obvious malicious indicators and appears to come from a familiar sender. This predisposes the recipient to trust the message and respond quickly.

When an attacker sends email from an actual compromised account, traditional headers offer little value. The email:

- Originates from legitimate infrastructure
- Passes authentication checks
- Appears indistinguishable from normal correspondence

This is precisely why reputation-based SEGs struggle with third-party ATO phishing. In these cases, detection must rely on something deeper than infrastructure or known threat indicators.

# How INKY identifies third-party account takeover

INKY detects third-party ATO phishing by building and maintaining sender profiles based on both technical and behavioral characteristics of email communication. These profiles capture what a sender's "normal" email looks like over time.

One of the most important behavioral techniques INKY uses is stylometry — the analysis of writing style to infer authorship. Originally developed to evaluate disputed texts, stylometry examines features such as vocabulary, sentence structure, punctuation, greetings and sign-offs to model how an individual typically writes.

INKY combines stylometric features with header-level signals to create a holistic representation of each sender it encounters. While it is often impossible to definitively prove that a message was not written by the claimed sender, INKY can determine when a message deviates significantly enough from established patterns to warrant caution.



**Caution:** External (nancyuser@comcast.net)  
Potential Sender Forgery [Details](#)

In the "Nancy" example, several subtle signals contributed to INKY's assessment:

- Unusual punctuation in the subject line
- Absence of a greeting or customary sign-off
- Structural anomalies, such as an empty To: field (often displayed as "undisclosed recipients")

Taken together, these deviations caused the message to stand out as an outlier compared to Nancy's normal communication style.

When INKY identifies this type of anomaly, it presents a yellow caution banner to the recipient, signaling that something about the message is unusual and encouraging the user to slow down, verify the request or confirm sensitive actions through a secondary channel.

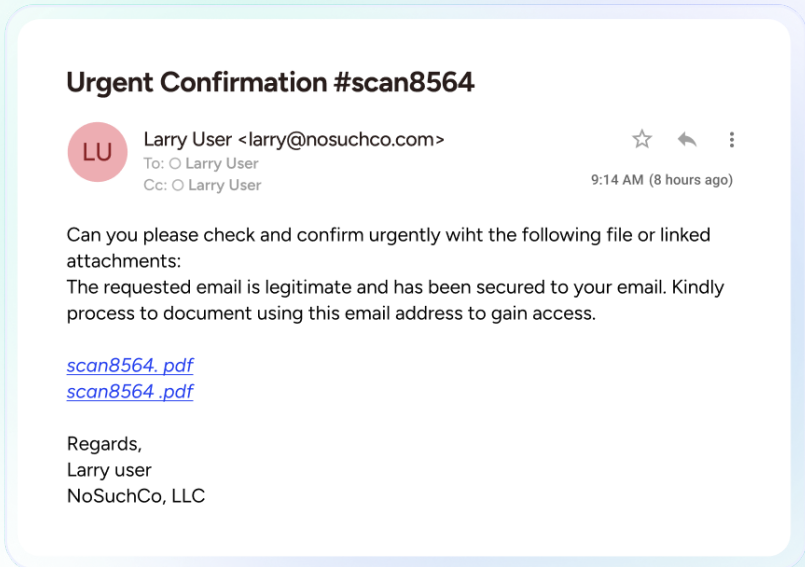
## A second example: Stylometry in action

The email below represents another third-party ATO attempt INKY detected.

This message exhibits many common characteristics of account takeover phishing: urgency, awkward phrasing, inconsistent punctuation and suspicious links that were not yet listed on threat intelligence feeds at the time of delivery.

Despite these signals, the email passed through the upstream secure email gateway undetected.


INKY flagged the message because its writing style differed substantially from the historical pattern associated with the sender, "Larry." Much like questioning the authorship of a disputed literary work, INKY applies stylometric analysis to assess whether a message plausibly matches the sender's established profile.




# How INKY uses clustering to detect ATO

Stylometry alone does not fully explain how INKY determines whether a message is likely legitimate or suspicious. To make these determinations at scale, INKY applies clustering, a machine learning technique that groups similar data points together based on shared characteristics.

Clustering enables models to identify patterns without requiring predefined labels. A common illustration is handwritten digit recognition, such as the MNIST dataset, where different representations of the same digit naturally group together based on visual similarity.



Each cluster represents examples that are more similar to one another than to those in other clusters. While visual examples are intuitive, the same concept applies to email — except that emails are represented in a much higher-dimensional space defined by dozens of extracted features.



In INKY's case, each email is mapped into a multi-dimensional feature space based on header attributes, stylometric traits, structural patterns and contextual indicators. Emails that appear to originate from the same sender typically cluster closely together. When a message falls far outside a sender's established cluster, it is treated as a potential impersonation or account takeover attempt.

The challenge — and the innovation — lies in selecting the right features and mappings so that similarity reflects how humans perceive legitimate communication. This process requires continuous experimentation and refinement, combining statistical rigor with practical insight into how real-world phishing attacks evolve.

INKY has spent years refining these models, continuously improving its ability to distinguish legitimate email from sophisticated impersonation — even when attackers use real accounts, real domains and real conversations.



# Why account takeover remains one of the most dangerous email threats

Account takeover continues to be one of the most pervasive and destructive phishing techniques because it exploits trust at its source. When attackers gain control of legitimate email accounts, they inherit reputation, relationships and credibility — allowing malicious messages to bypass technical defenses and manipulate recipients with alarming effectiveness.

Third-party account takeover attacks are particularly challenging. Organizations cannot prevent these compromises at the source, and traditional email security tools often lack the visibility needed to identify them once they occur. As attackers increasingly rely on real

accounts, real conversations and real context, detecting ATO requires understanding who normally sends an email, how they write and how their messages typically behave over time.

INKY applies advanced machine learning techniques, including stylometry and clustering, to build dynamic sender profiles and identify deviations that signal potential impersonation. By analyzing emails in context and focusing on sender authenticity rather than static indicators, INKY helps organizations detect and stop account takeover–driven phishing attacks before they result in fraud, data loss or business disruption.

## 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<sup>®</sup>

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](http://www.kaseya.com).

**[kaseya.com](http://kaseya.com)**

©2026 Kaseya Limited. All rights reserved. Kaseya and the Kaseya logo are among the trademarks or registered trademarks owned by or licensed to Kaseya Limited. All other marks are the property of their respective owners.