Kaseya®

WHITEPAPER

Understanding phishing: How a ransomware attack unfolds

Introduction

High-profile ransomware attacks impacting health care providers, global enterprises and shared technology platforms now dominate cybersecurity news. While ransomware is commonly associated with encrypted systems and ransom demands, few organizations understand how these attacks begin or why legitimate-looking emails so often play a central role.

In this understanding phishing guide, we trace the evolution of ransomware and explain how modern phishing techniques set the stage for attacks that progress from a single inbox to organization-wide disruption.

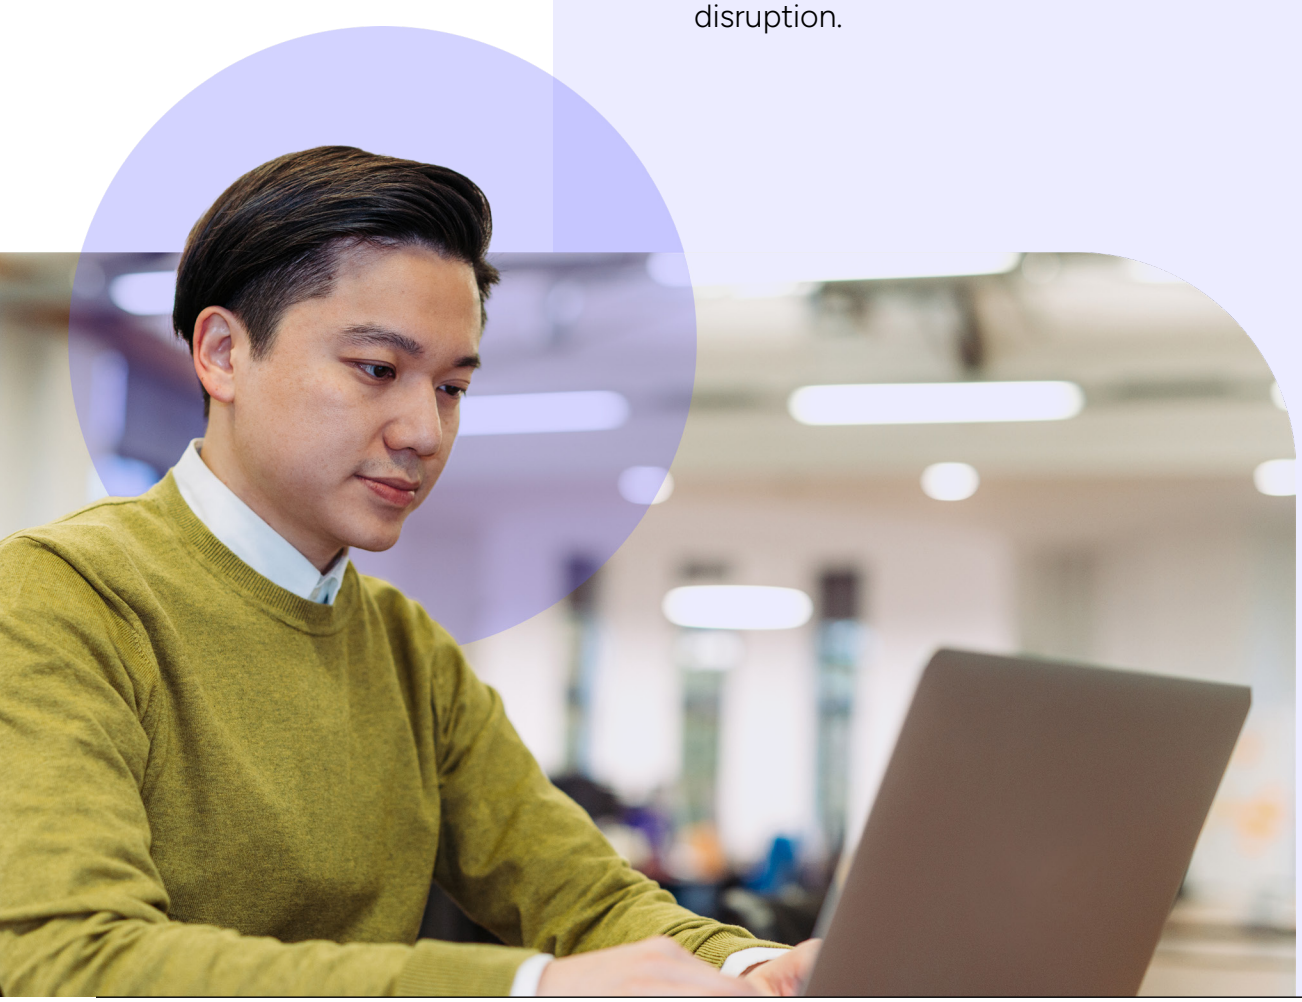

The origins of ransomware

Ransomware is not a new threat. Its origins date back more than three decades, long before today's cloud environments and AI-driven attacks. The first known ransomware incident, the AIDS Trojan, appeared in 1989 and was distributed via infected floppy disks mailed to conference attendees. Once installed, the malware encrypted filenames and demanded payment to restore access. While primitive by today's standards, it introduced the core ransomware model: deny access, demand payment.

For many years, ransomware remained a niche threat due to technical limitations and the difficulty of monetization. That changed in the mid-2000s as stronger encryption methods became widely available and digital payments became easier to transfer anonymously. By the early 2010s, ransomware had evolved into a scalable cybercrime model, setting the stage for widespread adoption by organized criminal groups.

The rise of Ransomware-as-a-Service (RaaS)

A major inflection point occurred around 2015–2016 with the emergence of Ransomware-as-a-Service (RaaS) kits. These platforms allowed skilled developers to package ransomware tools and lease them to affiliates, dramatically lowering the barrier to entry for attackers. Attackers no longer needed advanced technical skills — they simply needed a way in.

RaaS transformed ransomware into a mature criminal ecosystem. Affiliates handled initial access — often through phishing emails, credential theft or malicious attachments — while operators maintained the malware, encryption infrastructure and payment systems. Profits were split, incentives were aligned and attacks scaled rapidly.

High-profile outbreaks like Locky (2016) and WannaCry (2017) demonstrated how effective this model could be. Ransomware was no longer opportunistic malware — it was a business, complete with customer support portals, negotiation teams and branding. This shift accelerated both the volume and sophistication of attacks across industries.

From opportunistic attacks to strategic campaigns

Beginning around 2019–2020, ransomware underwent another fundamental transformation. Attackers moved beyond simple "encrypt and extort" tactics and adopted strategies designed to maximize pressure, impact and leverage.

Targeting organizations of all sizes

Ransomware evolved from mass consumer campaigns into a scalable model capable of targeting small businesses, large enterprises and entire supply chains alike. SMBs are targeted for ease of entry and faster payouts, while large enterprises are targeted for their data, scale and ability to disrupt entire ecosystems. Automation, phishing campaigns and AI-generated lures have made mass targeting both efficient and profitable.

Supply chain and platform-level attacks

Around 2020, attackers increasingly shifted toward supply chain targets — software platforms, service providers and shared infrastructure. Compromising a single vendor can create downstream access to thousands of organizations. Attacks involving widely used platforms and service providers such as Salesforce integrations, distribution partners like Ingram Micro, or enterprise platforms such as Oracle illustrate how ransomware groups now seek blast radius, not just individual victims.

Expanded outcomes beyond financial gain

Modern ransomware groups pursue outcomes far beyond ransom payments. Since approximately 2019, "double extortion" and "triple extortion" models have become standard. Attackers steal data before encryption and use it as leverage, threatening public exposure, regulatory penalties or operational disruption.

In critical sectors such as health care, ransomware has resulted in patient harm, delayed care, and safety risks. In global enterprises, attacks have caused supply chain disruptions, halted manufacturing and created geopolitical ripple effects. The objective is no longer just payment — it is maximum leverage, pressure and impact.

The modern ransomware era: 2024–2026

From 2024 onward, ransomware entered a new phase defined less by novel malware and more by precision, scale and abuse of trusted systems. While encryption remains a core tactic, today's campaigns focus on how attackers gain access and how much leverage they can create once inside.

Phishing has become more targeted, trusted and harder to detect

Between 2024 and 2026, phishing has evolved from obvious deception into high-trust impersonation, often abusing legitimate platforms and services. Attackers increasingly rely on:

- Legitimate email infrastructure and cloud services
- Authenticated senders and trusted domains
- Real links to real platforms with attacker-controlled content

Rather than spoofing domains, adversaries exploit platform trust, shifting malicious intent into file names, preview text and contextual cues that bypass traditional detection. This evolution reflects a broader shift away from technical evasion and toward behavioral manipulation at scale.

Initial access is more valuable than the ransomware itself

In this period, ransomware groups began treating initial access as a standalone commodity. Access brokers specialize in phishing-based credential theft and session hijacking, then sell that access to ransomware affiliates.

The result is faster, more reliable attacks:

- Phishing provides access
- Affiliates perform reconnaissance and lateral movement
- Ransomware deployment becomes the final step, not the first

This division of labor has accelerated ransomware timelines from weeks to days and reduced the likelihood of early detection.

Platform and identity abuse outpace malware innovation

Modern ransomware campaigns increasingly succeed without deploying malware at all in early stages. Instead, attackers abuse:

- Cloud collaboration tools
- Identity systems and OAuth permissions
- Shared SaaS environments and integrations

Because these activities occur within trusted platforms, they blend into normal business behavior. This has made identity compromise and email-based access more reliable than exploiting software vulnerabilities.

Impact is measured in business disruption, not just dollars

From 2024 onward, ransomware outcomes are evaluated by operational and systemic impact, not just ransom payments. Attacks increasingly aim to:

- Disrupt global supply chains
- Interrupt healthcare and critical services
- Create regulatory, legal and reputational consequences
- Apply pressure through data exposure rather than encryption alone

In many cases, encryption is optional. Stolen data, trusted access and the ability to disrupt business operations provide sufficient leverage.

Legacy email security models are structurally mismatched

These trends expose a fundamental gap in legacy email defenses. Systems optimized to block spoofed senders, known malware or bad domains struggle when:

- The sender is legitimate
- Authentication passes
- Links point to trusted platforms
- Malicious intent is embedded in context rather than code

As ransomware continues to evolve, email remains the primary ignition point, but the signals required to detect it have shifted from reputation-based indicators to behavioral, contextual and intent-based analysis.

Where ransomware still begins

As ransomware has evolved, one factor has remained consistent: email remains the most reliable and scalable entry point. Phishing enables attackers to bypass perimeter defenses, exploit human decision-making and initiate ransomware campaigns that unfold across organizations, supply chains and entire industries.

How attackers turn trusted platforms into phishing infrastructure

This campaign abuses legitimate Google Drive share notifications to deliver phishing and malware while leveraging the trust in Google's email infrastructure. The emails are genuinely sent by Google, pass authentication and contain real Google Drive links but the content is attacker-controlled. Threat actors use Unicode confusable (homoglyph) characters in the sender display name, subject, and file name to impersonate trusted entities or brands (In this case a known Law firm that the recipient uses and has done business with before) while evading keyword-based detection.

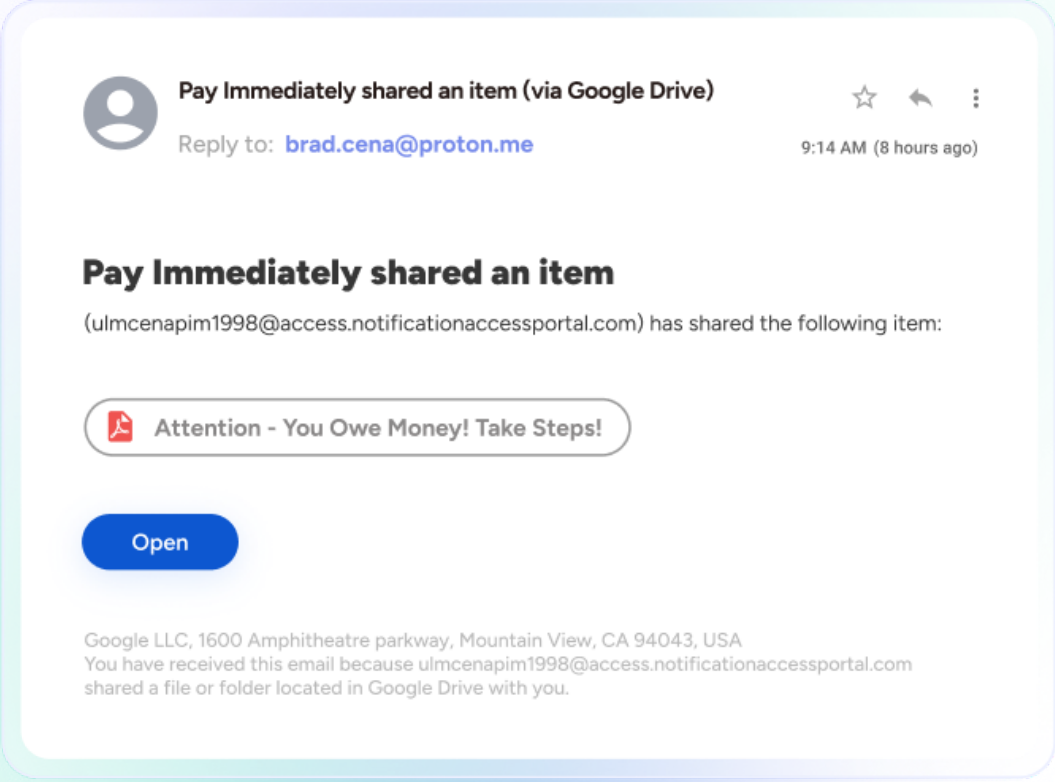

The social engineering payload is shifted into the file name and preview text, often using legal or financial urgency to prompt clicks. Victims are led to a real Google Drive URL hosting malicious documents, links or scripts, which further lowers suspicion because the domain is trusted.

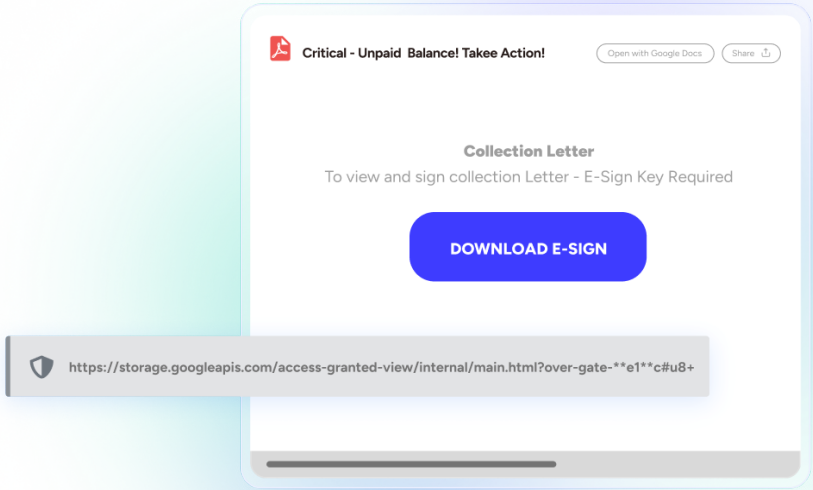

Additionally, the Reply-To address resolves to a suspicious, newly registered domain (notificationaccessportal.com) created within the same week the email was delivered, indicating it was purpose-built for this campaign; any reply to the message is routed not to the impersonated brand, but directly to the attacker's mailbox, as the bad actor registered the domain and used an associated email address to create a Google account and distribute the phishing via legitimate Google Drive notifications.

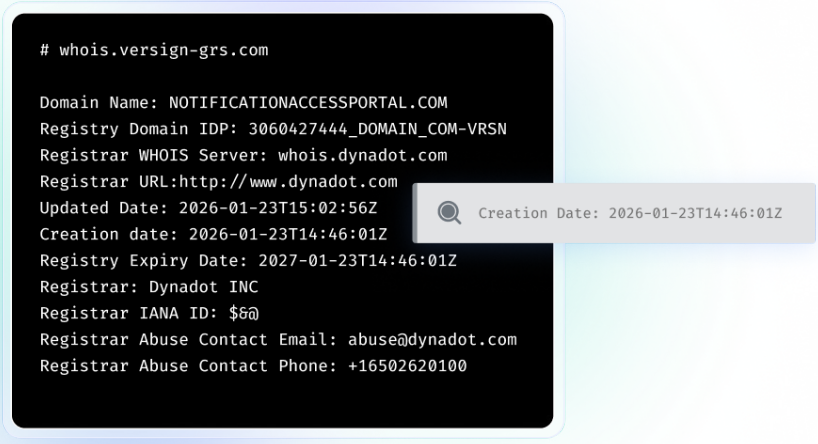

This technique is a form of platform abuse (or "Abused Service") phishing, not email spoofing. Traditional defenses struggle because transport and the sender's domain reputation is clean, and blocking Google Drive outright is impractical. Effective detection relies on behavioral and content signals such as mixed-script text, brand impersonation, first-time senders and high-risk themes rather than sender reputation alone.

How INKY identified the mail as phishing

INKY's AI was able to identify the malicious intent of the message by correlating multiple signals rather than relying on a single indicator. The platform detected the use of confusable Unicode text (reflected by the GenAI warning banner, in the dashboard screenshot), analyzed the surrounding context and correctly recognized that while the email was a legitimate Google Drive notification, it was being used as an abused service for phishing. In addition, INKY identified that the Reply-To domain was newly created, further increasing the overall risk score and reinforcing the malicious classification.

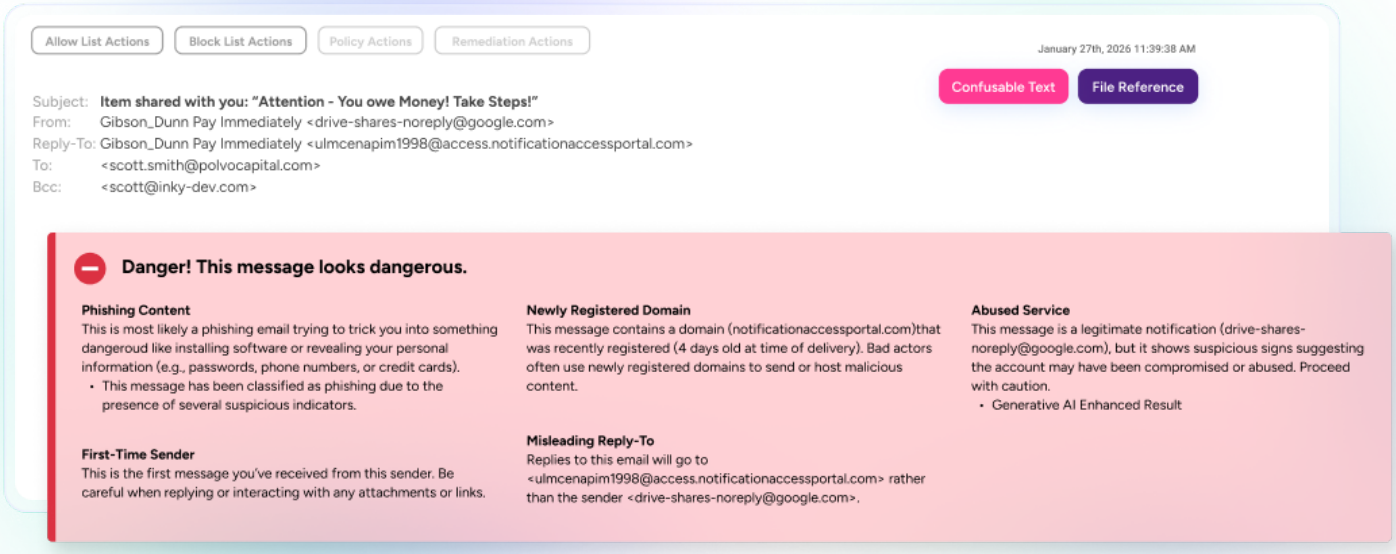

Colonial Pipeline kicks off opportunistic attacks

In this ransomware attempt — caught by INKY — the bad guys took advantage of the successful phishing-led ransomware attack executed against Colonial Pipeline to try to scare their targets into downloading malware. By posing as an internal sender (Help Desk), the attackers purported to being doing the opposite: requesting users to download an "update" that would protect against malware.

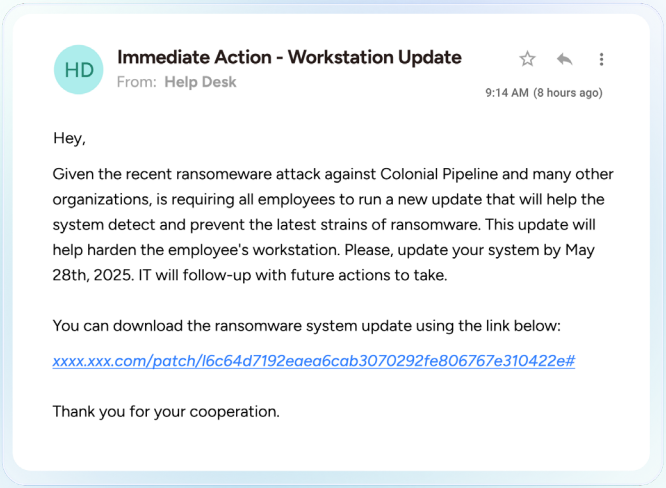

The black hats bought newly registered domains from NameCheap, and they immediately began sending phishing emails from those domains to targeted recipients.

NameCheap has, shall we say, "liberal" customer policies, for example, allowing all manner of registrations and accepting cryptocurrency as payment for domains and cloud hosting services. Given that ease of doing business, the attackers used the same domains to both send the phishing emails and host the malware. How convenient! In a nice touch, the bad guys used the target company's name as a subdomain in the malicious link to make it look like the email was an internal request. Here's the slick malware injection site hosted on ms-sysupdate.com:

Note that the target company's name used as a subdomain, but its logo and brand imagery were also sprinkled around liberally on the landing page. The big blue Download button initiated a download of a file called "Ransomware_Update.exe".

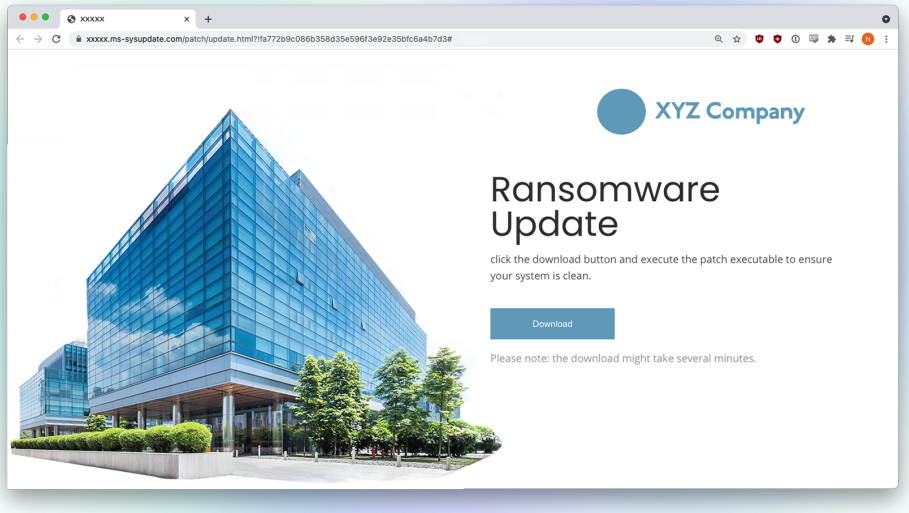

Payload: Cobalt Strike

Rather than protecting against malware, however, the payload was itself malware, in fact a particularly pernicious variety called Cobalt Strike. Although Cobalt Strike started out as a legitimate tool used for penetration testing, its source code was leaked in 2020. Since then, it has been ferociously abused by phishers.

Common antivirus (AV) systems, which focus on security data, often miss Cobalt Strike, which implements two main techniques to avoid AV detection: it obfuscates its own shellcode and leverages a domain-specific language called Malleable Command and Control (Malleable C2).

To detect executables, AV systems commonly implement sandboxing, which provides a separate environment to run and inspect suspicious executables. Cobalt Strike, however, hides shellcode over a named pipe. If the sandbox doesn't emulate named pipes, it will not find the malicious shellcode.

In addition, the attacker can modify and build their own techniques with the Cobalt Strike Artifact Kit.

After establishing itself in a system, Cobalt Strike can mimic popular services (e.g., Gmail, Bing, Pandora) to evade detection. The platform uses Malleable C2, which lets attackers modify Cobalt Strike command-and-control (C2) traffic at will. The attacker can then identify legitimate applications within the target organization, such as Amazon traffic, and modify the C2 traffic to appear as Amazon traffic using any number of publicly available profiles.

Cobalt Strike is highly customizable and can be used for command execution, key logging, file transfer, SOCKS proxying (to get around firewalls), privilege escalation, port scanning and lateral movement.

INKY prevents phishing that leads to ransomware shutdowns

As the reader will note, INKY caught the phish in the preceding examples. To prevent ransomware, organizations need to stop the attack at the beachhead. All it takes is one successful phishing exploit and the network falls to the mercy of the attacker. INKY is the most advanced email security solution on the market.

How INKY works

INKY takes a fundamentally different approach to email security, designed for a world where phishing emails often arrive through trusted platforms, authenticated senders and clean infrastructure.

Rather than relying on static rules or sender reputation alone, INKY uses advanced Generative AI and complementary detection models to analyze the intent of every email in real time. Each message is evaluated inline, as it flows into the inbox, without delaying delivery or disrupting mail flow.

Intent-based detection powered by GenAI

INKY's GenAI analyzes the full context of an email, including language, structure, visual elements, links and sender behavior, to determine what the message is trying to make the recipient do. This concept-based analysis allows INKY to detect modern phishing techniques that evade traditional controls, such as:

- Abused trusted platforms and cloud services
- Brand impersonation using real domains and real links
- AI-generated social engineering and urgency tactics
- Messages that pass authentication but carry malicious intent

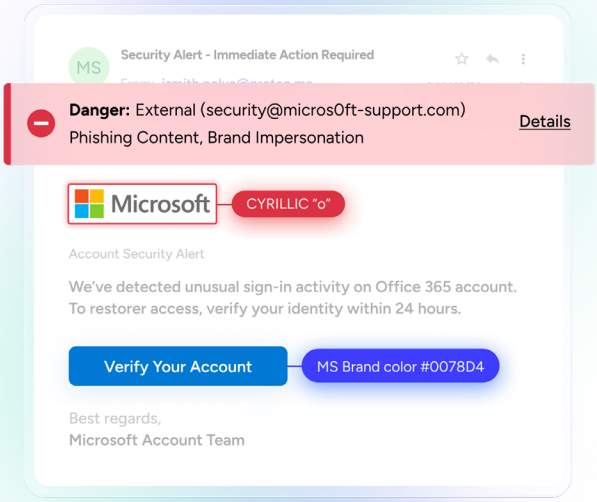

Visual and brand forgery detection

To combat sophisticated impersonation attacks, INKY performs computer vision analysis in a secure browser environment. Logos, branding elements and layout patterns are extracted and compared against known brand assets and the actual sending domain.

This enables INKY to detect brand forgery and lookalike attacks even when no malicious payload is present and the sender infrastructure appears legitimate.

Behavioral learning and identity awareness

INKY continuously learns how people communicate inside and outside the organization. Shortly after deployment, it builds behavioral sender profiles based on real communication patterns, including:

- Writing style and formatting habits
- Typical signoffs and tone
- Devices and email platforms used
- Geographic and behavioral consistency

When an incoming email deviates meaningfully from a known sender's profile, INKY flags the message as a potential impersonation attempt and clearly alerts the recipient.

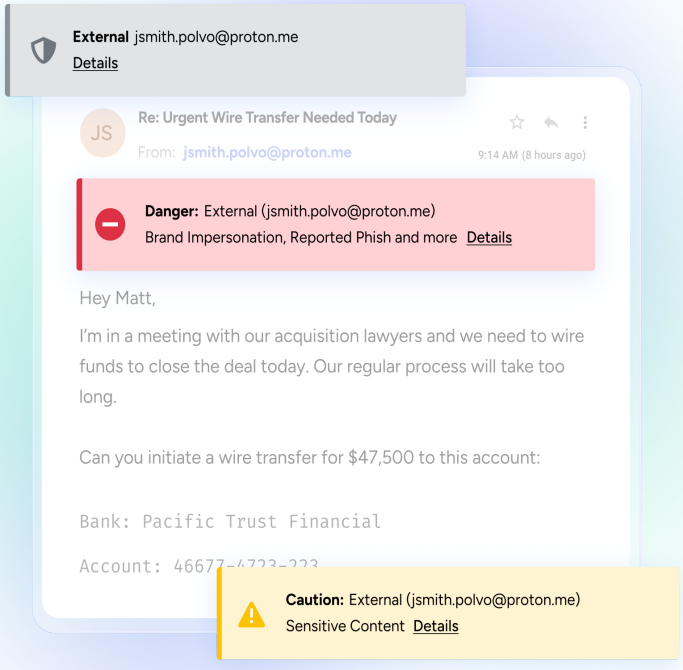

All analyses run in parallel and complete in seconds. Instead of issuing a binary allow-or-block decision, INKY assigns a risk assessment that is communicated directly to the user through the Email Assistant with a clear, color-coded banner.

Real-time user coaching at the moment of risk

INKY doesn't just block threats; it guides users at the exact moment a decision is required. When an email triggers risk signals, the banner explains why the message is suspicious, turning every flagged email into a teachable moment.

This real-time coaching has been shown to significantly improve user decision-making and reduce repeat risky behavior, helping organizations lower email-related incidents without relying on periodic training alone.

Reducing risk before ransomware can begin

Modern ransomware attacks rarely start with malware. They start with a trusted email and a single human decision.

INKY prevents ransomware by stopping phishing where it actually begins, in trusted email interactions that legacy tools were never designed to see.

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®

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