

State of Phishing Report




TABLE OF CONTENTS

- 01** Executive Summary / Introduction
- 02** Key Findings
- 03** Overall Email Threat Trends
- 06** Threat Category Analysis
- 11** Live Scanning and Link-Based Phishing
- 12** Non-email Attack Vectors
- 13** Next Steps



Executive Summary

The current AI-powered threat landscape reveals an unprecedented surge in attack volume, with a 202% increase in phishing messages in a 6-month period, and credential phishing attacks rising 703% in the same period. Organizations face a barrage of advanced attacks on the mailbox, while mobile and collaborative applications are growing rapidly as threat vectors. The public success of [ShinyHunters and Scattered Spider](#) highlight the critical shift from email-only to multi-channel attack vectors.

Our analysis also shows attackers are able to spawn multiple versions of an attack to overwhelm users and security analysts. To this effect, we've seen 80% of malicious links in attacks are previously unknown zero-day threats, demonstrating that traditional threat intelligence and signature-based detection methods are increasingly ineffective against modern, AI-powered attack campaigns. The threat of today will not be the threat of tomorrow.

Looking ahead to 2026, we expect this rapid evolution to accelerate, with AI-generated attacks becoming more sophisticated and harder to detect, while attackers increasingly target messaging platforms beyond email, including business collaboration tools, SMS, and social media. The bottom line is email isn't the only human risk factor; there is a broader messaging security problem that requires a fundamental shift in how organizations approach threat detection and prevention.



Introduction

Welcome to the first edition of Varonis' **State of Phishing Report**, where we uncover the most critical insights from the current threat landscape. We analyze the key attacks and trends over the past year or more, using data to determine whether email, messaging, or collaboration security threats are escalating, stabilizing, or declining. Discover what to expect in 2026 and the emerging attack vectors you must watch out for.



Key Findings

Before diving into the detailed analysis, let's examine the critical insights that emerged from our research. These key findings highlight the most significant trends and developments that security leaders should consider when evaluating their threat prevention strategy.



EMAIL ATTACK VOLUME

Overall email attack volumes increased 202% over a 6-month period, and are trending up as we approach 2026.



ADVANCED PHISHING FREQUENCY

Users receive at least **1 advanced phishing link every week** that bypasses traditional network security or gateway controls.



MOBILE THREATS

Users encounter up to **600 mobile threats annually** on average.



TOP EMAIL ATTACK VECTOR

Malicious email link-based attacks are the top email attack vectors, with email text-based attacks following close behind.



NEW AND UNKNOWN THREATS

On average, **80%** of the links in email link-based threats are **new and unknown**.



SOCIAL ENGINEERING

Social engineering continues to be ever-present, **rising 141% over six months**.



CREDENTIAL PHISHING UP

Breaches often start with attackers logging in. It's no surprise, then, **that credential phishing is up 703%**.

Overall Email Threat Trends

Let's look at the overall trends for email-born attacks over a 12-month period. When you look at these graphs, please note all figures presented represent advanced threats detected per 1,000 mailboxes and delivered on a weekly basis (unless otherwise noted). We measure the prevalence of threats per 1,000 mailboxes to provide a clear, consistent view of risk across businesses of all sizes. Notably, these advanced threats bypass legacy and first-generation AI email security systems.




FIGURE 1.

Overall advanced threat index covering all email-based threat types

+	+	+
+	+	+
+	+	+
+	+	+

ADVANCED EMAIL THREAT INDEX




Focusing on the latter six month period, we saw a steady and consistent rise in the overall volume of phishing attacks. As you can see in Figure 2, mapping the changing nature of the volume of phishing attacks and adding an overall trend line shows this consistent rise.




FIGURE 2.

Six month overall email threat trend and trendline

SIX MONTH ADVANCED EMAIL THREAT INDEX





202%

**INCREASE IN THE
NUMBER OF PHISHING
MESSAGES THROUGH
THE USE OF AI**


Within this six month snapshot, the number of attacks per 1,000 mailboxes each week has increased linearly. Currently, we are capturing close to one advanced attack per mailbox each week. As we reach the 1,000 thresholds, this translates to nearly one advanced attack for every single mailbox each month. This steady increase indicates a substantial volume problem that individual technologies or manual efforts cannot handle effectively.

Unfortunately, this rise in the volume of attacks isn't a huge surprise. Throughout the year, we've shown evidence of attackers using AI-designed campaigns to evade detection, automate their processes, and target victims at scale. The use of AI has led to a remarkable 202% increase in the number of phishing messages delivered per 1,000 mailboxes.

This trend underscores a significant shift in email security dynamics. We are now operating in what can be described as a "volume game," where the sheer number of attacks overwhelms traditional security measures. Relying solely on Security Operations Center (SOC) analysts to manually manage this influx is proving unsustainable and may set organizations up for failure.

FIGURE 3.
Email threat categories (text, link, and file-based) trends over time

THREAT INDEX PRIMARY CATEGORIES



**NO SINGLE THREAT
CATEGORY
CONSISTENTLY
DOMINATES.**

- + **Text-based threats** (no attached payload)
- + **Link-based threats** (malicious links as payload)
- + **File-based threats** (file attachment as payload)



While text-based threats have gained wide attention since the FBI began tracking them in 2013, link-based threats remain the primary challenge in terms of volume. Text-based threats like Business Email Compromise (BEC) attacks, a subset of Social Engineering attacks, are important attack vectors that should be addressed. At the same time, the data in Figure 3 shows that link-based phishing is the primary vector, and text-based threats are second. The data shows these attack methods alternate throughout the year in prevalence — **no single threat category consistently dominates**. This reinforces why browser-based threat defense is a critical and necessary part of a phishing defense strategy.



LINK-BASED THREATS REMAIN THE PRIMARY CHALLENGE IN TERMS OF VOLUME

File-based threats remain constant, though they’re generally less prevalent than other attack vectors in our context.

This reduced frequency can be attributed to several factors:

- + Organizations implementing restrictive file-based policies
- + Enhanced email security infrastructure
- + Secure file transfer mechanisms
- + Improved operational procedures

Modern file-based threats have evolved beyond traditional self-executing files.

They now often incorporate techniques like HTML smuggling and don’t always contain the malicious payload themselves but instead side-load that content once the file is executed. These more advanced file attacks focus on:


- + Credential harvesting
- + Personal Identifiable Information (PII) theft

Threat Category Analysis

The current threat landscape shows several prominent attack vectors, with some key trends emerging:

FIGURE 4.

Specific threat type focusing on credential phishing and social engineering



Recent data shows a significant rise in both credential phishing and link-based threats. Shown in Figure 4, **credential phishing is up 703%**. These attack methods frequently overlap, as many credential phishing attempts incorporate malicious links as part of their strategy. However, while there's substantial overlap, they're not entirely the same. Recent data indicates a late surge in these tactics as we move towards the holiday season. This was associated heavily with the major phishing campaign abusing DocuSign APIs. **Social Engineering attacks are up 141%**, maintaining a consistent presence throughout the year and showing steady growth.

[BLOG](#)

MIND GAMES: HOW SOCIAL ENGINEERING TACTICS HAVE EVOLVED

[Read the article](#)

Credential Phishing

Credential harvesting attacks begin when a malicious email bypasses security filters and reaches a user’s inbox. The attacker’s primary goal is to lead victims to a webpage where they will enter their login credentials. Let’s use a recent DocuSign campaign as a sample.

FIGURE 5.

Legitimate emails sent from trusted vendors leveraged by threat actors

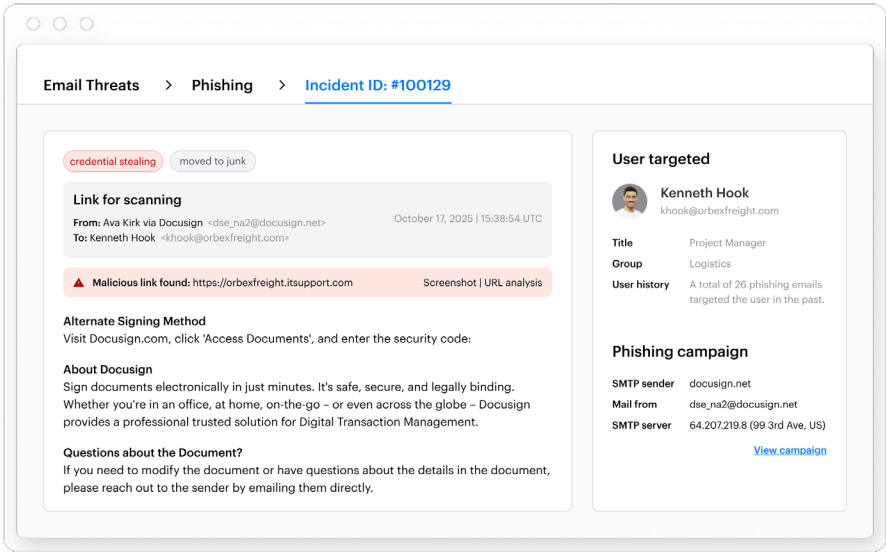



Figure 5 shows an email that successfully passed through email filters because it originated legitimately from DocuSign, a trusted business service. The email can’t be blocked without impeding normal business operations, as employees need to exchange DocuSign documents with customers. The security challenge arises because while the email legitimately uses DocuSign’s cloud infrastructure, it contains a direct link within its environment that leads to malicious content.

There are three main methods attackers use to direct users to these malicious pages:





Taking the simplest case of a malicious link, the attack flow includes multiple layers of obfuscation:

First, attackers employ redirectors and intermediate sites to mask the destination. They often implement testing mechanisms to filter traffic—either based on specific criteria like browser age and source or sometimes random filtering for high-volume campaigns. This helps keep their phishing pages active longer by making them harder for security services to detect.

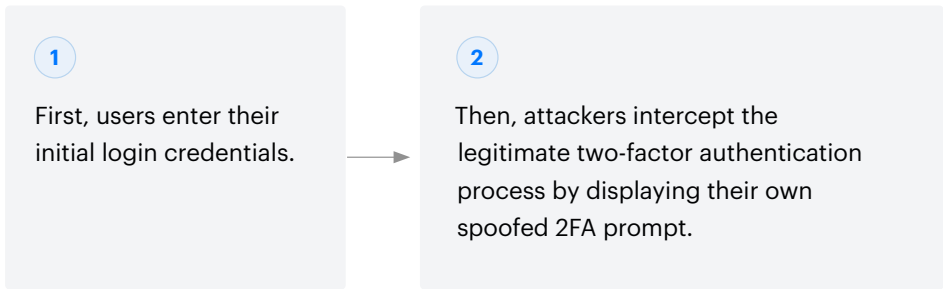
When users reach the destination, they typically encounter a “human verification” test.

This might appear as:

- + CloudFlare turnstile services
- + Google CAPTCHA pages

After passing these verification steps, users face the actual phishing attack: a spoofed login page designed to capture credentials.

This is a two stage process:



Once attackers have both the initial credentials and the two-factor authentication details, they can successfully compromise the account.

Social Engineering

Social engineering continues to be ever-present, as shown in Figure 4, rising 141% from the previous six months.

Social engineering encompasses several attack styles, including but not limited to:

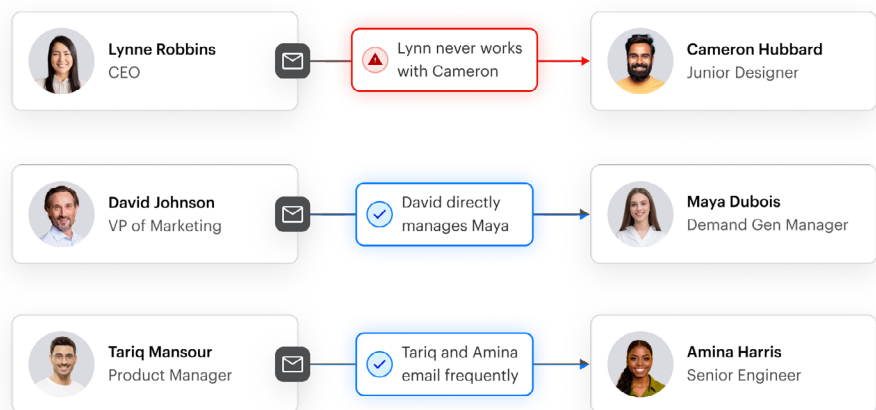
- + Quote requests for vendor partnerships
- + Fake loan and purchase scams
- + Business Email Compromise (BEC)

There are also BEC subcategories within Social Engineering that include:

- + Payroll theft
- + Reconnaissance
- + Invoice fraud

While BEC attacks receive significant market attention and can cause substantial financial damage, they represent a small percentage of the total attack volume. Organizations have improved at detecting traditional BEC, leading attackers to change tactics frequently as we noted previously.

The most dangerous attacks now involve compromised external accounts, where recipients are accustomed to communications from the sender. Without social graph analysis to understand relationship patterns, these attacks are particularly effective. We capture these by leveraging machine learning, neural networks, or generative AI.





The most dangerous attacks now involve compromised external accounts.

A notable development is the ShinyHunters campaign, which introduced a new multi-channel approach:

- + Flooding inboxes or initiating phone calls from someone posing as IT support
- + Instructing the victim to install a modified version of Salesforce's Data Loader
- + Tricking the employee to enter their 8-character device code as a legitimate request
- + Gaining standing access to Salesforce data or performing actions within the platform

Social engineering attacks have increased significantly year over year, with trends showing:


- + Integration with other phishing techniques
- + More sophisticated multi-dimensional approaches
- + Combination of social engineering, link-based threats, and multi-channel phishing
- + Some variants include malicious files for endpoint compromise

The trend indicates continued steady growth with increasing complexity rather than one-dimensional attacks.

BLOG

AI-GENERATED PHISHING:
HOW ONE EMAIL TRIGGERED A
GLOBAL NPM SUPPLY CHAIN CRISIS

Read the article



Live Scanning and Link-Based Phishing

The image below illustrates the detection methods for link-based threats, currently the leading category of email-borne attacks. The graph shows two categories:


- + Known threats detected through existing threat intelligence (blue)
- + Threats detected through real-time scanning of previously unseen URLs (red)

The primary takeaway is that, on average, 80% of the links in email link-based threats are unknown. These pages are created moments before being sent, last for a very short period, and are de-weaponized just as fast. In other words, most malicious links will slip past controls relying on known threat feeds. We utilize a proprietary virtual browser to analyze and live scan URLs when analyzing email threats.

FIGURE 6.

Comparison of malicious link identification: live scan (unknown) vs. threat intel (known bad)

LINK-BASED DETECTION METHODOLOGY



Throughout the year, the data reveals that relying on threat intelligence alone is ineffective against modern attacks. Most malicious links sent in email are zero-day URLs, created and deployed moments before the attack using AI and machine learning tools available on the dark web. These tools can generate thousands of unique pages through neural networks and phishing kits.

The key finding from this data is that traditional signature-based detection methods, which require 24–48 hours to develop signatures for malicious URLs, are no longer sufficient. The graph demonstrates why real-time scanning capabilities are essential for detecting these zero-day threats as they're being deployed.

80% OF THE LINKS IN EMAIL LINK-BASED THREATS ARE UNKNOWN



Non-email Attack Vectors

URL and Browser Threats


The data shows links that bypassed traditional network security layers and DNS-based detection mechanisms that users clicked on but were mitigated by Varonis using its AI-embedded browser security. This graph in Figure 7 is key to understanding the bigger trend: **phishing attacks have evolved beyond email-only delivery channels.**

FIGURE 7.

Advanced browser threats getting through advanced network security reaching users through multiple alternative channels such as:

- + Browser-based messaging apps
- + Teams/Slack communications
- + LinkedIn messages
- + Personal Gmail accounts
- + CRM forms

ADVANCED BROWSER THREATS



Notable examples in the data include operations from noteworthy groups, such as Lazarus and Midnight Blizzard. Even organizations with multiple network layers and browser protections are vulnerable because of these attacks:

- + Utilize brand-new, previously unseen threats
- + Bypass categorization due to their novelty
- + Can be hosted on legitimate, trusted infrastructure (notably DocuSign, OneDrive, and SharePoint)
- + Pierce through organizational controls before detection systems can update

Users get advanced phishing every week.

We take this all-encompassing approach because spot protection for individual applications isn't scalable. Varonis customers have found browser-based phishing protection to be a crucial addition in mitigating multi-channel attacks. Each link shown in the graph successfully circumvented existing security controls, highlighting how attackers are adapting their delivery methods as email security hardens.

Next Steps

As the human risk landscape continues to evolve, so must your defense strategy. AI-powered hackers are flooding users with social engineering attacks across email and collaboration tools like WhatsApp, Slack, Teams, and Zoom. By imitating tone, branding, voice, and even video, they’re launching automated campaigns that humans — and traditional email security tools — can’t detect. Worse, these attacks are overwhelmingly targeted at data and lead to breaches.

Varonis email security solutions are powered by predictive AI models that use techniques such as computer vision, natural language processing, and virtual browsers to protect users from the widest range of social engineering attacks. Extending our data-centric threat detection capabilities with the market’s best phishing and social engineering detection solution, we are able to stop the most prevalent forms of attack and extend our end-to-end approach to protecting data.

If you have any questions about the data in this report or how to prevent these advanced attacks, please visit us at varonis.com. To discover how many threats may be bypassing your current defenses, try our Email Risk Assessment. It takes 5 minutes to set up and provides valuable insights to help you plan your defenses.



Get started with a free Email Risk Assessment.

Reduce your risk without taking any. Contact our team to learn what will be covered in your **free** Email Risk Assessment.

Get started at varonis.com/platform/email-security

ABOUT VARONIS

Varonis (Nasdaq: VRNS) is the leader in data security, fighting a different battle than conventional cybersecurity companies. Our cloud-native Data Security Platform continuously discovers and classifies critical data, removes exposures, and detects advanced threats with AI-powered automation.

Thousands of organizations worldwide trust Varonis to defend their data wherever it lives — across SaaS, IaaS, and hybrid cloud environments. Customers use Varonis to automate a wide range of security outcomes, including data security posture management (DSPM), data classification, data access governance (DAG), data detection and response (DDR), data loss prevention (DLP), and insider risk management.

Varonis protects data first, not last. Learn more at www.varonis.com.