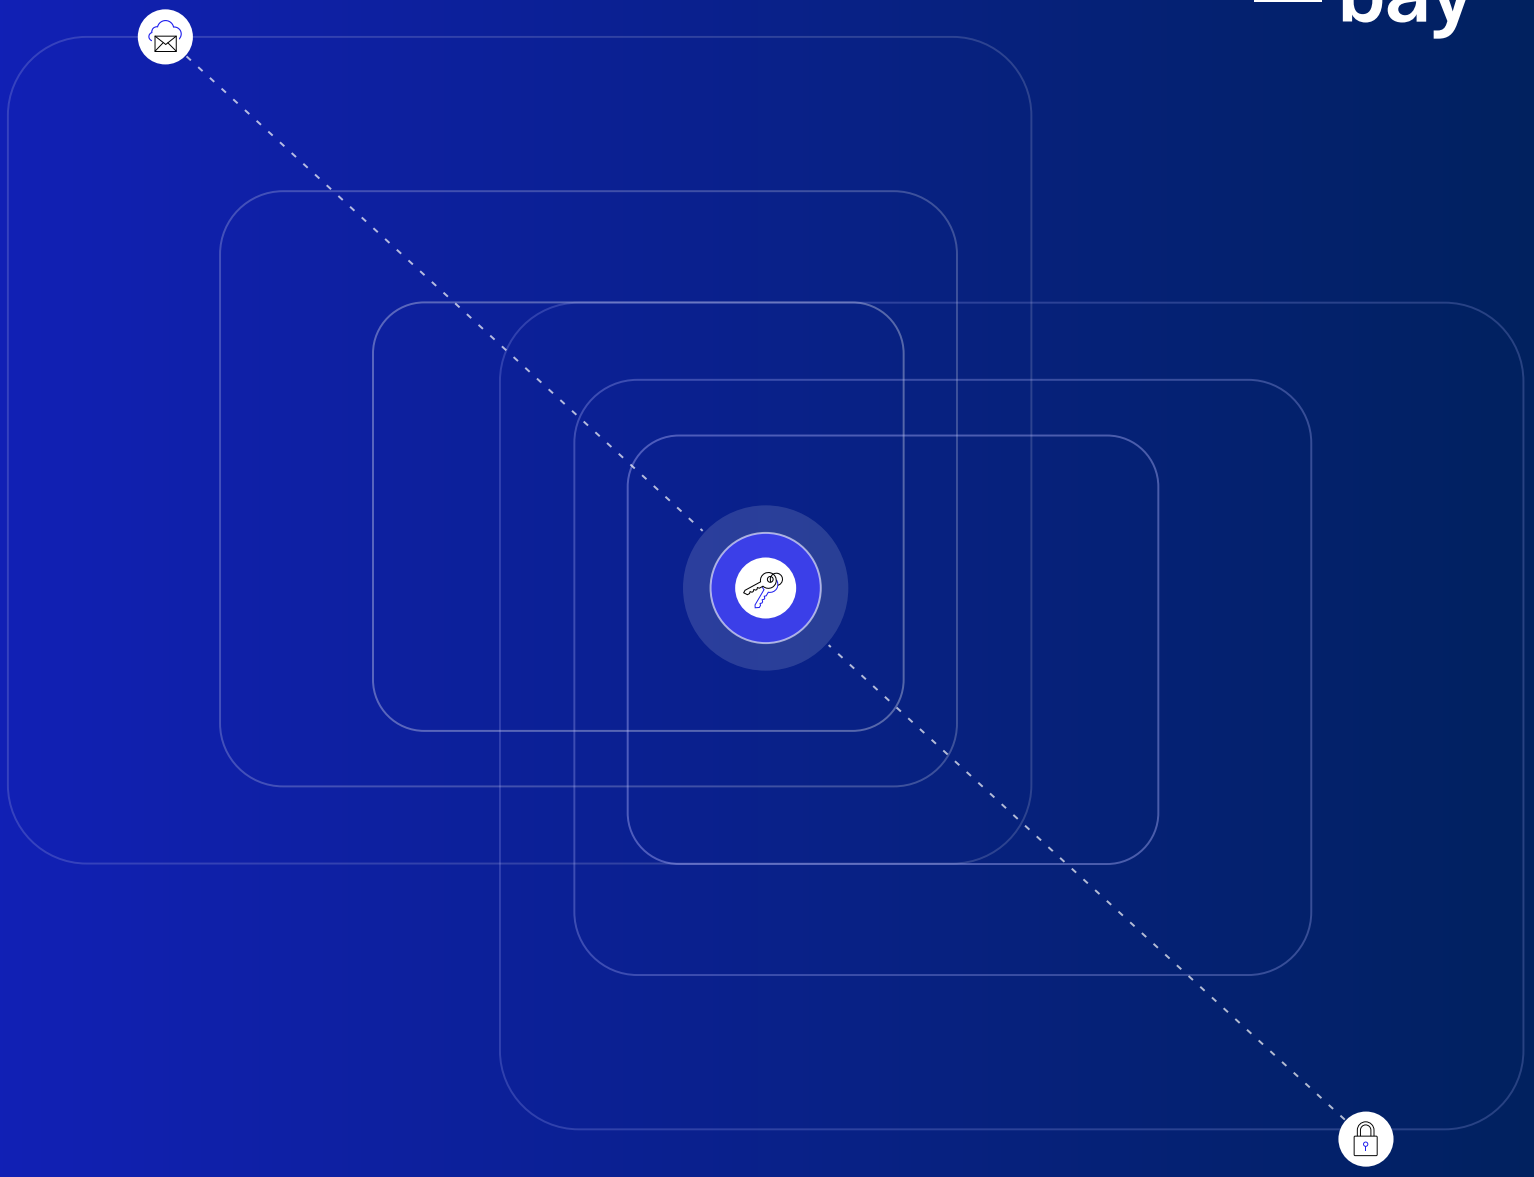
at bay



Email +
Remote Access

2026 InsurSec Rankings Report

An analysis of At-Bay claims + cybercrime data

Table of Contents

Introduction	4
Chapter 1: Email Rankings	6
Deep Dive: Brand Impersonation and Abused Infrastructure	15
Chapter 2: Remote Access and VPN Rankings	21
Conclusion	27
Methodology	28
Contributors	30

Introduction

Every year, At-Bay's InsurSec Rankings Report asks a simple question: What are the risks associated with the decisions businesses make about the email and remote access technology they use? In past editions, we answered that question for email and remote access separately. They were the two most common ways attackers got in, and each had its own rankings and its own advice.

This year, we're continuing to do that, in addition to illustrating how the two connect through credential harvesting and abuse. The pattern shows up across our claims, our MDR data, and our incident investigations, and demonstrates why investing in security on both fronts is critical.

On the email side, this year brings slightly better news. In last year's report, attackers appeared to be using AI faster and more effectively than defenders, and detection slipped across nearly every product we track. For 2025 through Q1 2026, that trend flattened, as new AI capabilities helped email security providers regain ground. Email claims continue to be elevated, however, ever since 2024 when generative AI first became widely available.

In a new analysis from At-Bay Security, 10% of emails flagged by MDR for Email relied on brand impersonation or abused infrastructure to slip past existing email security. Credential theft was the leading intent behind those messages, and those stolen credentials can then be used to log in through remote access tools like VPN and RDP, where many ransomware attacks begin.

That's where the news is worse. Having any VPN in place still carries more risk than having none, with some remote access tools being correlated with even higher risk.

Customers using SonicWall saw their risk climb from 5.8X to 8.4X year-over-year.

That increase was driven largely by Akira ransomware, and in most cases Akira got in by logging in with legitimate, stolen credentials rather than by exploiting a new vulnerability. In At-Bay Response & Recovery investigations where the initial access vector was known, VPN and RDP accounted for 65% of ransomware cases. In extortion cases where threat actors explained how they got in, nearly 60% said the credentials were either harvested or purchased on the dark web.

Stopping credential theft entirely may not be realistic. What businesses can do is make sure a stolen credential isn't enough on its own. That means enforcing phishing-resistant MFA everywhere, having detection that someone actively monitors, and moving away from architectures that grant access to the entire network on a single login. Yet 60% of the businesses in At-Bay's Post-Cyber Event Hardening (PCEH) program still had MFA missing or only partially deployed, even after they had suffered an attack.

As our InsurSec Report showed earlier this year, cyber outcomes are not randomly distributed. The tools a business chooses, and how it configures them, change what happens next. This report shows where each tool stands, and how attackers move between them, to help businesses make those choices well.

Key Findings

1

AI-driven email defenses closed the detection gap

Email claims have remained elevated since 2024, but this year's data shows lower email claims frequency for policyholders using Proofpoint, Mimecast, and Barracuda, the top three Secure Email Gateways included in our analysis, each of which added new AI-based detection.

2

Customers with Google Workspace continue to see elevated email claims frequency

Organizations using Google Workspace saw the lowest email claims frequency for the fourth year in a row, but continued to see elevated claims frequency compared to findings in our 2024 report. Meanwhile, Microsoft 365 users continued to see improved outcomes YoY.

3

Proofpoint and Mimecast reclaimed the top spots for secure email gateways

The average claims frequency of At-Bay customers with secure email gateways dropped 17% year-over-year, reclaiming some of the ground they lost in previous years. Proofpoint (down 26%) and Mimecast (down 17%) were the only email security solutions with claims frequency below average. Both have added major AI-based detection since 2024.

4

Brand impersonation and abused infrastructure drove 10% of emails flagged by At-Bay MDR

Of the messages At-Bay's MDR for Email flagged between February and August 2026, 10% were brand impersonation or abused infrastructure attempts. Credential theft is the leading intent behind these attacks.

5

Ransomware risk associated with SonicWall VPNs climbed from 5.8x to 8.4x year-over-year

Businesses with on-premise VPNs continued to see elevated risk of ransomware, with customers using SonicWall seeing the highest relative risk. Businesses with any VPN continued to see higher risk than those with no VPN at all.

6

60% of policyholders who experienced a claim had MFA gaps

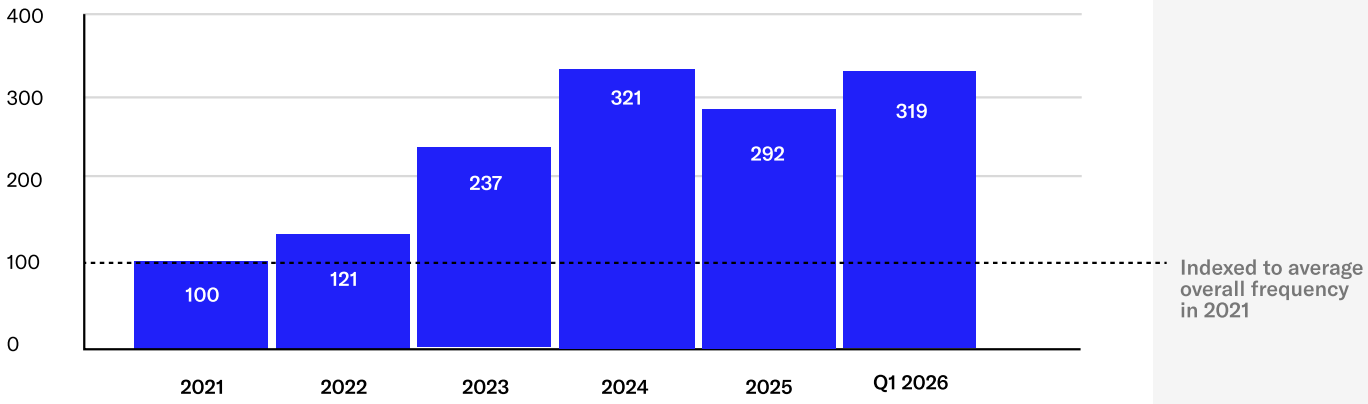
When At-Bay's Security team helped policyholders strengthen their defenses after a claim, 60% either hadn't deployed MFA or had it on only some systems. In a separate At-Bay analysis, 30% of policyholders with claims over the past four years believed they had MFA when they didn't.

Email Rankings

FIGURE 1

Email Claims Frequency Remains Elevated at About 3X the Levels of 2021

Indexed Email Frequency by Year



In 2024, email claims frequency saw an increase of 35% year-over-year and has remained elevated at about 3X the levels of 2021 ever since. We attribute these increased levels of email claims to the emergence of AI and attackers' aggressive investment and effectiveness at identifying ways to accelerate and improve their processes through the use of AI (we cover this in depth in our recent 2026 Guide to AI Risk).

In the 2025 InsurSec Rankings Report, we observed a noticeable decline in apparent malicious email detection capability for two major email solution providers (Google and Microsoft) and most email security products in 2024 and Q1 2025. Our hypothesis was that, since many of the email security tools we track had only begun introducing significant AI-based capabilities in the latter half of 2024[1], we would not be able to see their impact until some time in 2025. We anticipated that, in that period, we would see newly integrated AI capabilities in email security tools take back some of the advantage that attackers seemed to have gained in earlier years.

Those capabilities have arrived, but attackers have adapted as well. The story of email security for the period covered by this year's report (2025 to Q1 2026) is the story of an artificial intelligence arms race, with neither attackers nor defenders definitively gaining the upper hand.

[1]Sources: (1) Mimecast, Mimecast Fortifies Email and Insider Risk Defenses with AI-Powered Enhancements, October 2024; (2) Proofpoint, Proofpoint Sets New Industry Standard in Email Security with Adaptive Threat Protection Capabilities Across the Entire Email Delivery Chain, May 2024

Email Solutions Rankings

Google and Microsoft are the two most prominent email providers tracked by At-Bay. While they aren't email security companies per se, the native security features built into their respective platforms play an important role in establishing a baseline of security capability for companies that rely on them for email service. And, for many organizations that have opted not to purchase a discrete email security solution, the security provided by the email solution itself will be the only email security they have.

FIGURE 2

Google Workspace Customers Saw the Best Outcomes for the Fourth Year

Email Solutions Rankings by Claims Frequency

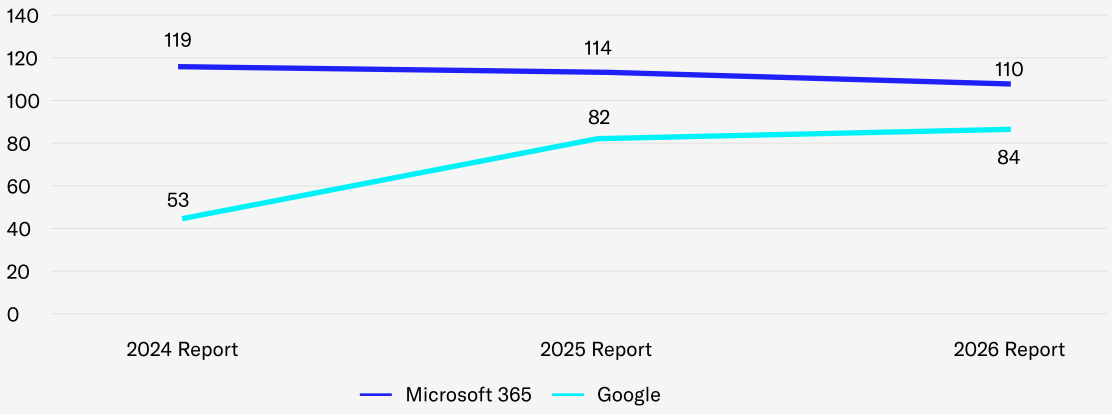
Email Solution	Indexed Email Claims Frequency (2025 - Q1 2026)	% Change YoY
Google	.155%	+2%
Microsoft 365	.185%	-12%
Average Frequency for All	.176%	-9%

For that reason, we have historically tracked the performance of Google and Microsoft side-by-side with market-leading email security companies as though they were direct competitors. Both vendors made meaningful upgrades to their native email security capabilities during the period covered by this report. However, while Google Workspace customers have seen the best outcomes since we began publishing this analysis, they've seen a significant jump in email claims frequency since our 2024 report, including this year. Businesses with Microsoft 365, on the other hand, have continued to see a decrease in email claims frequency.

FIGURE 3

Microsoft 365 Continues to Close the Gap with Google

Indexed Email Claims Frequency, YoY (Microsoft and Google)



Gmail Security Improvements



Data Loss Protection

Google added additional features to strengthen its email platform against account takeovers and the use of email to leak sensitive information. In February 2025, Gmail Data Loss Protection (DLP) became generally available[2]. While not specifically targeted at email-based fraud, the unauthorized use of email accounts to send fraud messages might be detected more easily if messages contain information that is tracked by the DLP function (e.g., bank account numbers, routing numbers, etc.).



AI Threat Protection

In April 2025, Google announced that it had added a new holistic AI threat protection model to Gmail[3]. This feature acts as a supervisor over existing AI, machine learning (ML), and rule-based defenses and reviews the signals those systems produce. By combining everything Google can see across sender behavior, message content, and other signals, it makes a more accurate judgment about whether something is malicious. This approach is akin to the techniques that a skilled human analyst might use to weigh several small, uncertain clues together to determine whether the full picture represents a threat or just something out of the ordinary.



Account Compromise Protection

May 2025 saw the addition of device-bound session credentials and improved audit and investigation telemetry for Gmail[4]. These two features reduce the risk of account compromise via certain types of replay attacks and improve the ability of security teams to perform analyses of account takeovers, respectively.

[2]Source: Google Workspace Updates, Workspace data loss protection (DLP) for Gmail is now generally available, February 2025
[3]Source: Google Workspace, New in Gmail: Making end-to-end encrypted emails easy to use for all organizations, April 2025
[4]Source: Google Workspace Updates, Gmail will soon stop support for the 3DES encryption cipher for incoming SMTP connections, May 2025

Microsoft 365 Security Improvements

LLM-Based BEC Detection

Microsoft's most important recent email security enhancement arrived in November 2024[5] when they introduced LLM-based detection for Business Email Compromise (BEC)[6]. As At-Bay has reported in the last several editions of our InsurSec Report, our most common claim type is financial fraud, and malicious emails are the most common starting point for the majority of claims generally. Thus, an AI-based feature to detect instances of compromise for an individual email account or for the email tenant as a whole represents a significant opportunity to identify the starting point for a potential fraud campaign before it has a chance to cause damage.

Mail-Bombing Detection

2025 email security enhancements for Microsoft 365 included the introduction of a feature to detect mail-bombing[7] — a tactic wherein an attacker floods a target with legitimate subscription and newsletter traffic. This tactic is sometimes used to distract defenders from other, more serious activity that can be a precursor to fraud or a computer intrusion. Its addition to Microsoft 365 provides a meaningful enhancement to companies' ability to stop this class of email threats and may reduce the occurrence of follow-up incidents such as ransomware attacks.

Improved Security Feedback and Remediation

Microsoft improved the security feedback and remediation loop for its email products in 2025 by enabling defenders to dispute Microsoft's assessments about the maliciousness of email addresses and URLs and to share with Microsoft any allow / block list actions taken at the tenant level8. The latter improvement potentially enables a significant enhancement in Microsoft 365's awareness of the current email threat landscape by leveraging the collective security wisdom of its customers.

At-Bay Stance Fraud Defense

Stance™ Fraud Defense helps protect against costly financial fraud by adding a layer of email security that catches sophisticated fraud attempts that legacy email security solutions miss. Earlier this year, we launched mailbox takeover detection for users with a Microsoft 365 integration, automatically detecting malicious email rules that attackers use to hide their activity — catching compromised mailboxes before they're weaponized for financial fraud or data theft.

Learn more at at-bay.com/stance/fraud-defense/

[5]As with similar features introduced by Mimecast and Proofpoint around the same time, this was too late for the 2025 edition of At-Bay's InsurSec Rankings Report to have reflected its impact.
[6]Source: Microsoft Ignite, What's new in Microsoft Defender for Office 365, November 2024
[7]Source: Microsoft Ignite, What's new in Microsoft Defender for Office 365, June 2025
[8]Ibid.

Email Security Rankings

In At-Bay's 2025 InsurSec Rankings Report, we noted that Sophos appeared to have taken the lead in effectiveness over previous leaders Proofpoint and Mimecast. Our hypothesis at that time was that Sophos had been an early mover in implementing a particular type of detective capability that was well-suited to that moment's email threat landscape: Natural Language Processing (NLP).

This year's report reflects the full impact of the significant AI capabilities introduced by other vendors from late 2024 through the first half of 2026. The collective impact of these new capabilities have eroded the apparent advantage gained by Sophos. This year's top-performing email security vendors are again our historical top performers from 2024, with Proofpoint in the top spot and Mimecast (2024's winner) following close behind.

FIGURE 4

Proofpoint and Mimecast Return to the Top of the Rankings

Secure Email Gateway Ranking by Claims Frequency

Secure Email Gateway Provider	Indexed Email Claims Frequency (2025 - Q1 2026)	% Change YoY
Proofpoint	.079%	-26%
Mimecast	.090%	-17%
Barracuda	.143%	-7%
AppRiver	.162%	-19%
Intermedia	.166%	-10%
Sophos	.196%	+133%
Average Frequency for all At-Bay Policyholders With Secure Email Gateways	.109%	-17%

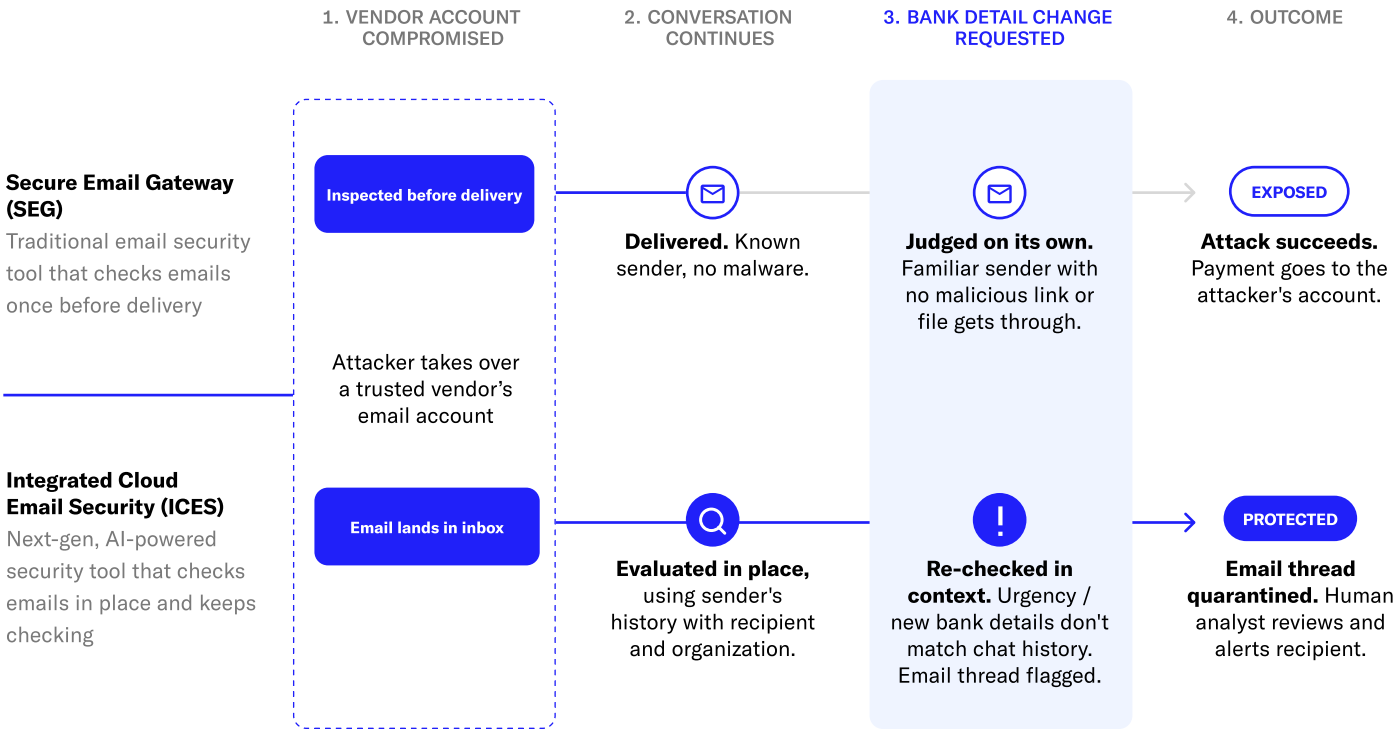
A Shift in Approach for Email Security

When At-Bay began tracking email security products in 2021, the dominant model was the Secure Email Gateway (SEG), which inspects emails in line, before they reach the recipient's server. The biggest trend At-Bay observed this past year was SEG vendors adding API integration alongside that in-line processing, so they can now reassess email after delivery.

Even so, SEGs still evaluate each email largely on its own, with little weight on context like the history between a sender and recipient. That leaves a gap for attacks from known, trusted senders, such as a compromised vendor account. Meanwhile, many organizations have moved from on-premise email servers to cloud-based solutions, largely because their APIs connect email to business tools like CRM systems, code repositories, and ticketing platforms.

Those APIs enabled a new type of tool: Integrated Cloud Email Security (ICES), which evaluates emails in place after they land in the inbox, usually before the recipient is notified. From inside the mailbox, ICES tools build context that gateways traditionally lack, like the recipient's and organization's history with a sender. Replies and new messages trigger re-inspection, and if a sender turns out to be compromised, the tool can remove related messages organization-wide.

SEGs have closed some of the gap with ICES tools, but conversational context is still the clearest difference. Consider a compromised vendor account. After a long history of legitimate conversations about payments, a message arrives from the same address urgently asking to change bank details for an ACH transfer. A SEG sees a known sender and no malware. An ICES tool sees language and banking details that don't match the conversation history, and flags the email thread, removing it from the inbox while it's reviewed by a human analyst who can assess and notify the recipient if a threat is identified.



Proofpoint Innovations

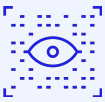


ICES Deployment

Proofpoint's biggest shift for the period covered by this report was architectural. In Q1 2025, Proofpoint launched Core Email Protection API which enabled customers to deploy their solution in an ICES configuration rather than requiring customers to route their email through a Proofpoint server (i.e., a SEG configuration)[9]. Proofpoint's announcement of this new deployment option described it as combining existing threat-protection engines with the AI email security capabilities it gained via its acquisition of Tessian in late 2023.

Tessian's contribution included behavioral and contextual techniques that are particularly useful against payloadless social engineering and anomalous communications — just the type of prolific email-borne threats currently afflicting At-Bay policyholders. Additionally, deployment as an ICES tool gained Proofpoint the ability to monitor the email tenant itself for threats such as account takeovers and tenant-level compromise, significant threats for which SEG products are wholly ineffective.

Closing the ICES Blind Spot



While At-Bay considers ICES tools to be a more modern and flexible approach to email security, they do have a key shortcoming we alluded to earlier. Specifically, inbound emails are not inspected until they arrive in a recipient's inbox.

Most ICES tools are fast enough to scan newly arrived emails before a recipient has a chance to interact with them, but this isn't always the case. Acknowledging this weakness, Proofpoint followed their 2025 product changes by launching a unified SEG plus API architecture, which became available in June 2026. This hybrid approach effectively allows customers to take advantage of both deployment models: pre-delivery inspection via a traditional SEG and post-delivery monitoring via an API connection.

For Microsoft 365 customers, this option also includes monitoring for internal email (i.e., company-internal emails that don't leave the tenant for delivery) and direct send paths, a feature that enables devices or applications to send email directly to a mailbox without passing through an email server. These last two updates further improve the ability to rapidly detect compromised accounts which mitigates the risk of email-borne fraud.

[9]Source: Proofpoint, Q1 Business Momentum Cements Proofpoint's Position as Undisputed Leader in Human-Centric Security, April 2025
[10]Source: Proofpoint, A new chapter for email protection: unifying SEG and API security, March 2026

Mimecast Innovations

Mimecast's 17% drop in email claims frequency most likely comes from the AI-powered BEC protection it introduced in October 2024[11]. In last year's report, we predicted that new AI-based detection from vendors like Mimecast would bring claims down among their customers. Those capabilities were in place for the whole period this report covers, and our data now reflects the improvement we expected.

Mimecast has kept investing since then. The two developments below arrived at the very end of our reporting period or after it closed, so they probably had little effect on the results we report here, but they're significant innovations that we expect to show up in next year's data.



Multi-Vector Threat Protection: Mimecast's standout enhancement for the period covered by this report was Multi-Vector Threat Protection (MVTP) which it began rolling out to select customers in December 2025 before making it generally available in January 2026[12]. Similar to Google's holistic AI "supervisor," MVTP correlates weak indicators across areas such as sender and domain reputation, spam signals, URLs, email content and, where available, customer social graphs[13]. This approach, which is similar to the way a human analyst might evaluate the maliciousness of an email that a recipient had reported as suspicious, is intended to identify threats where no single indicator provides definitive evidence on its own. In these instances, multiple correlated indicators might paint a clear picture of maliciousness, and that's the insight that MVTP is intended to surface.

Shortly after rollout, Mimecast asserted that MVTP had identified three million additional phishing messages across 20% of its customer base while maintaining a false positive rate below 0.002%[14]. This reporting, coming in March 2026, likely reflects the results observed among early adopters and so probably under-counts the actual impact of this functionality after it is deployed to their full customer base. We should note that At-Bay doesn't typically publicize or rely on vendor data for our reports. But, the 17% drop in email-related claims experienced in 2026 by At-Bay policyholders with Mimecast provides strong evidence that MVTP is delivering as promised.



ICES Deployment: Mimecast also made architectural changes to enable deployment in an ICES configuration while continuing to offer its legacy SEG deployment option. In March 2026, the full protection stack available from Mimecast became available via API. Customers now have the option to deploy the tool as a SEG, as an ICES, or both. However, unlike Proofpoint, Mimecast emphasizes that the same detective capabilities are available in both the SEG and the ICES configurations. So, while customers could deploy both simultaneously, it doesn't appear that Mimecast has intentionally structured the two deployment options to complement each other.

[11]Source: Mimecast, Mimecast Fortifies Email and Insider Risk Defenses with AI-Powered Enhancements, October 2024
[12]Source: Mimecase Community, Introducing Multi-Vector Threat Protection, December 2025
[13]Social graphs track interactions between email senders and recipients. An email between a sender and recipient who haven't previously exchanged emails is considered more suspicious than one between parties who communicate often.
[14]Source: Mimecast, Faster detection, fewer threats, zero compromise, March 2026

Other Notable Developments

The other email security vendors that are prominent among At-Bay policyholders, Barracuda and Sophos, made noteworthy improvements to their products from H2 2025 - H1 2026, but their impact wasn't enough to displace the historical leaders.

Barracuda

AI-Powered Threat Analysis: In May 2025, Barracuda joined the trend of leveraging AI to correlate and analyze multiple threat indicators to improve the accuracy of detection[15]. This enhancement is likely reflected in their improved performance among At-Bay policyholders for the period covered by this report.

ICES Deployment: In June 2026, Barracuda released Integrated Email Protection, their ICES deployment option[16]. This feature works with Google and Microsoft email systems through API integration. It keeps checking delivered messages over time, looking for signs of abuse and pulling back any message that turns out to be malicious after it's already landed in an inbox. This ensures that a message isn't simply forgotten after a one-time safety check. If a link is weaponized after delivery, or the sender's account gets compromised later on, Barracuda can still catch it and act. Among the email security vendors At-Bay tracks, this kind of ongoing, cross-domain risk checking appears to be distinctive. However, its impact isn't reflected in our current claims data and likely won't become clear until mid-2027.

Sophos

BEC and Impersonation Protection: In May 2025, Sophos substantially expanded BEC and impersonation protection with specific emphasis on external VIPs (important customers, vendors, and partners)[17]. With email-borne financial fraud continuing to be the most common incident type afflicting At-Bay policyholders and with most fraud incidents involving impersonation of third parties, this capability enhancement directly targets a top area of concern for At-Bay customers.

Warning, Detection, and Recall Features: 2026 Sophos updates included the addition of new warning banners based on mailbox-specific communication history[18], newly registered domain protection[19], trace and recall for internal messages that were determined to be malicious after an initial recipient forwarded them within the enterprise[20], and AI-based email header analysis[21]. This last feature appears similar to the capability offered by At-Bay's own Fraud Defense solution in leveraging email authentication data to determine maliciousness and also supports investigations by analyzing email headers to capture routing information.

Note that risk improvements from product updates made by Sophos in 2026 may not be reflected in the rankings for this report, as the data in this report reflects 2025 - Q1 2026. We'll know more about how 2026 updates impact the rankings in next year's report.

At-Bay Stance Cyber Advisory Services



Stance Advisory Services provides access to a team of At-Bay security experts who can help prioritize risks, resolve issues, and design better security controls. This includes an email threat assessment service, which evaluates the effectiveness of current email security capabilities, identifies ways to improve email security, and helps implement those improvements.

Learn more about At-Bay Stance Advisory Services →

[15]Source: Barracuda, Barracuda Unveils Next-Generation Threat Detection Powered by Multimodal AI, May 2025
[16]Source: Barracuda, Barracuda Unveils Integrated Email Protection, An Evolution in Email Security Delivering Cyber Resilience Made Easy for the Agentic AI Era, June 2026
[17]Source: Sophos Community, BEC (Account Compromise and VIP) Protection Enhancements, May 2025
[18]Source: Sophos Community, New Sender Banner (Address Historian), January 2026
[19]Source: Sophos Community, Newly Registered Domain (NRD) Protection, April 2026
[20]Source: Sophos Community, Contact Trace and Recall (CTR): Post-Delivery Protection for Internally Forwarded Emails, April 2026
[21]Source: Sophos Community, AI Analysis of Email Headers, May 2026

Deep Dive: Brand Impersonation and Abused Infrastructure

The email rankings above show how email security vendors are evolving to catch modern threats, but there's another side of that equation: how attackers are evolving to get a malicious email past those defenses and in front of an employee in the first place. Two tactics account for a sizeable amount of that effort in the data we track:

1. **Brand impersonation** is an email built to look like it came from a company the recipient already knows and trusts (e.g. Microsoft, DocuSign, PayPal) to convince the recipient they're dealing with the real brand.

2. **Abused infrastructure** uses legitimate systems (a freshly registered domain, an open redirect, or a platform like Meta or Monday.com) to deliver or host a message that can slip past email filters.

Both approaches continue to evolve alongside the detection capabilities meant to catch them, and understanding how they're used today, who they target, and what attackers actually do with what they gain matters well beyond a single flagged inbox. A single set of stolen credentials can be the difference between a blocked attempt and a full-scale ransomware incident.

At-Bay's MDR for Email provides 24/7 inbox protection powered by AI and monitored by dedicated security experts who review and respond to flagged messages. Of all flagged messages from February 1 to August 31, 2026 — over 2.2 million — 10% were either brand impersonation or abused infrastructure attempts. Those emails reached more than 72K individual inboxes across our book, and brand impersonation made up 61% of the email volume while abused infrastructure made up 39%. In the next sections, we'll do a deep dive on each of these tactics based on real world examples our MDR for Email is catching.

FIGURE 5

10% of Flagged Emails Are Abused Infrastructure or Brand Impersonation

Share of All Flagged Emails by Attacker Tactic

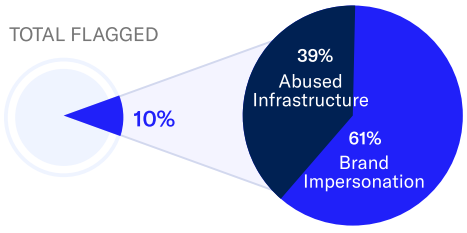
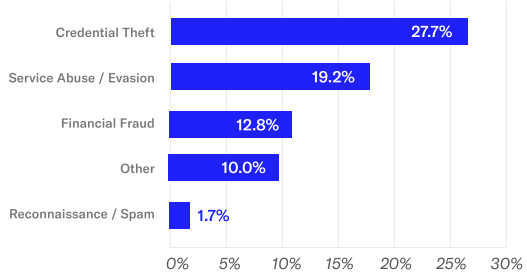


FIGURE 6

Credential Theft Leads Every Tracked Attacker Intent

Share of All Flagged Emails by Primary Intent (Excluding Unclassified)



Both tactics share the same underlying goal: borrowing an existing brand's credibility to get past a target's defenses. Credential theft is the single largest intent behind these emails, driving just over a quarter of the categorized messages we flagged as brand impersonation or abused infrastructure — more than financial fraud, service abuse, or any other tactic we tracked.

Credential theft refers to an email built to get the recipient to hand over a login, usually through a fake sign-in page or a login-gated prompt dressed up as a shared file or document request. It's the most direct route to what the attacker is actually after. A phished credential lets an attacker walk past impersonation entirely and log in as the real employee or vendor, which is why credential theft so often leads to account takeover, fraudulent instructions sent from a genuinely compromised inbox, or access to whatever other systems share that same login.

Which Brands and Services Attackers Borrow Most

Brand impersonation and abused infrastructure don't distribute evenly across the brands and services attackers lean on. A handful of names account for a disproportionate share of the volume in each category. The rankings below show which brands are impersonated most often and which services are most commonly abused.

FIGURE 7

Top 10 Impersonated Brands by Email Volume

Share of Brand Impersonation Flagged by At-Bay MDR for Email

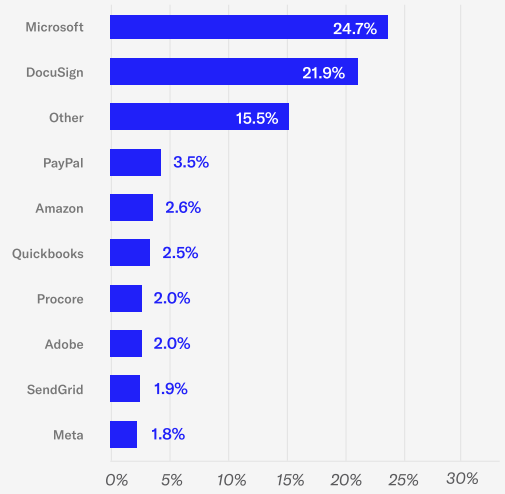
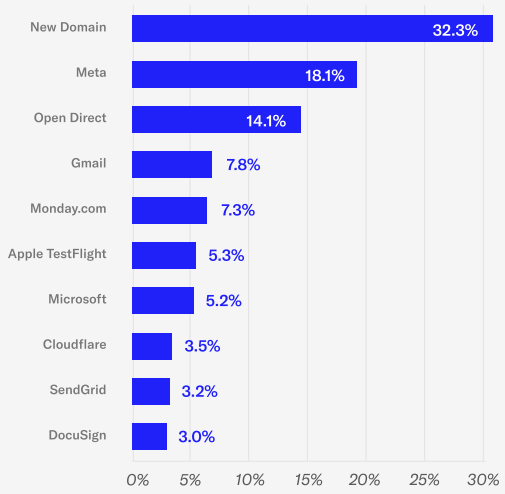


FIGURE 8

Top 10 Abused Infrastructure by Email Volume

Share of Abused Infrastructure Flagged by At-Bay MDR for Email (Excluding 'Other')



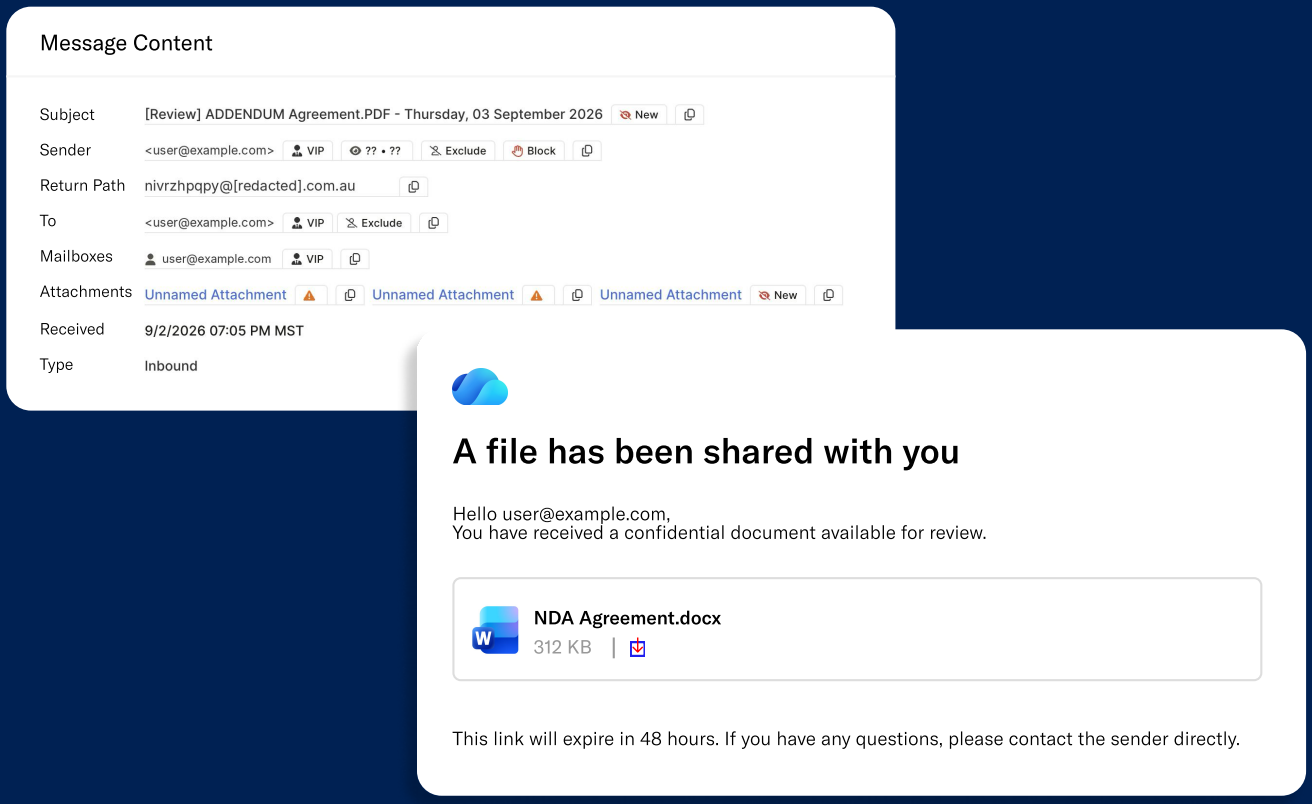
Microsoft and DocuSign alone account for 46.6% of brand impersonation volume, and the rest of the list reads like a typical company's software stack: PayPal, Amazon, QuickBooks, Procore, and Adobe. Attackers pick these brands because employees deal with them every day at work. A DocuSign signature request or a Microsoft file-share notification is a normal part of the workday, so one more in the inbox doesn't raise a red flag.

Newly registered domains top the abused infrastructure list at 32.3%. Email security tools judge a sender partly on its track record, and a domain that's only days old doesn't have one. With no history of abuse to flag, a new domain can slip past filters that would catch a known bad sender. Legitimate platforms like Meta, Gmail, and Monday.com work for the opposite reason: They come with years of trusted history that a malicious message inherits. The Meta example later in this section shows how that plays out step by step.

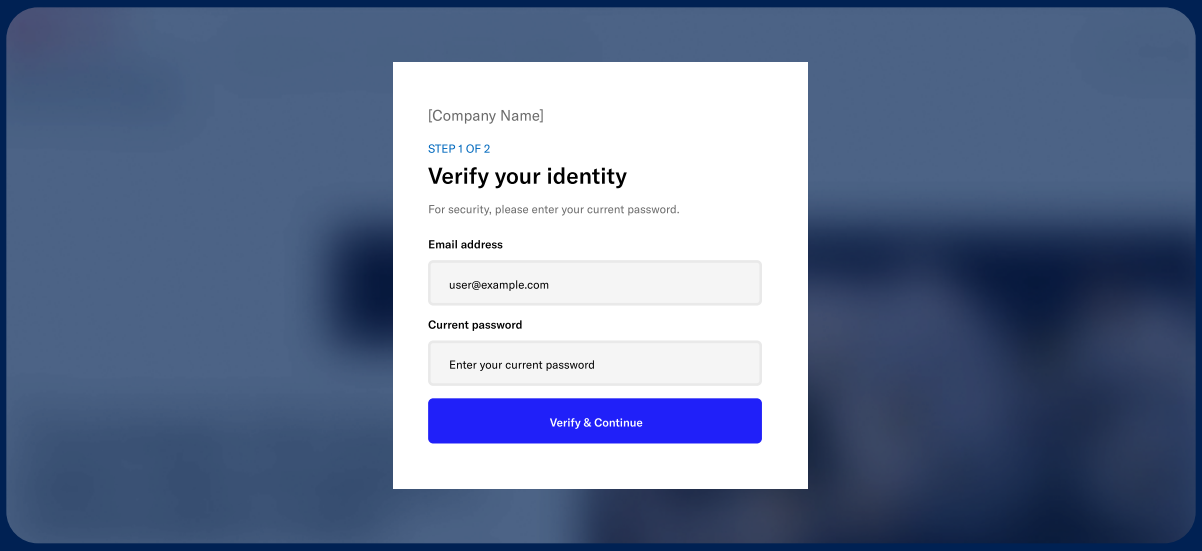
Anatomy of an Attack: Microsoft Brand Impersonation

At-Bay Security's MDR team observed this instance of Microsoft brand impersonation, which follows a typical progression for this type of attack.

Step 1: The targeted employee received an email that, when opened, appears to be from Microsoft, prompting them to open a link to view a document. The sender is spoofed to match the recipient's email address. The return path is the threat actor's actual email address, which would be the address to which any further messages were sent if the recipient replied to the message.



Step 2: The “NDA Agreement.docx” link opens the URL 'hxxps://aWHM1P6f-yPaL--.shoireditch[.]com/user@example.com', which leads to this spoofed page that appears to be a legitimate Microsoft login page. The page prompts the user to enter their Microsoft credentials before they can access the document.



Step 3: The employee inputs their login credentials, believing they’re using a legitimate Microsoft login portal, but it’s actually a cloned portal that sends credentials straight to the attacker’s infrastructure.

Why this works: Because the attacker faked the identity of a trusted company the recipient expects to hear from at work, they’re likely to believe it’s legitimate and input their login credentials. From there, the attacker can access whatever files, SharePoint sites, and email history the employee has legitimate access to, set up mail forwarding rules to quietly monitor conversations about invoices or payments, or send further messages from the employee's actual, trusted address to coworkers, vendors, or customers who would have no reason to doubt them.

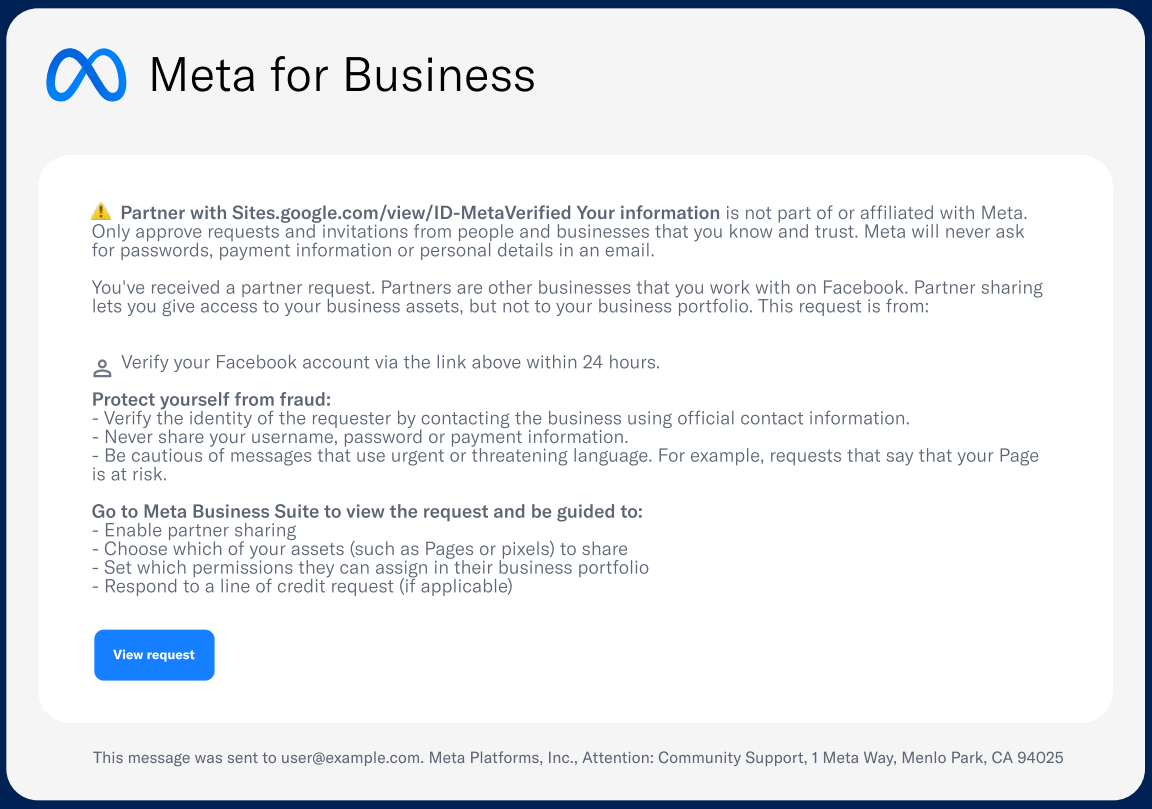
How to prevent it:

- **Train employees to check domain, not design.** The fake login page was a convincing copy of Microsoft's real one, but the link led to a domain with no connection to Microsoft at all. Training employees to check the address bar before entering a password, especially one reached through a document link, could stop this attack at step two.
- **Require phishing-resistant MFA.** A password and a text-message code can both be captured on a cloned page. A passkey or FIDO2 security key can't, because it's tied to the real Microsoft domain and won't authenticate against a lookalike, even if an employee clicks through and tries to log in.
- **Have employees watch for new inbox rules.** Once an attacker has valid credentials, setting up a quiet forwarding rule is a common next move. Alerts on newly created or modified mail rules can catch the takeover before it turns into fraud against a vendor or customer.
- **Train employees for the pattern, not just the brand.** This attack used Microsoft's name, but the trick underneath — an urgent document link that demands a login — works with any trusted brand. Awareness training built around real examples like this one keeps the pattern recognizable no matter which company gets impersonated.

Anatomy of an Attack: Abused Meta Infrastructure

At-Bay Security's MDR team observed this instance of abused infrastructure. It used two legitimate platforms, Meta and Google, to move a target step by step toward a site the attacker controlled.

Step 1: The threat actor set up a fake business account on Meta and sent a business partner request to a real business. The actor named the fake business "Sites.google.com/view/ID-MetaVerified" and filled the requester field with "Verify your Facebook account via the link above within 24 hours." That way, the notification reads like an instruction from Meta. Meta's own system generated and sent the email, so it was genuine Meta mail and passed standard authentication checks.



Step 2: Clicking "View request" opens a legitimate Meta partnership page that shows the requesting company. Nothing on this page is malicious. It's the real Meta site.

Step 3: On that page, the fake company's name shows up as a hyperlink. A target has to click through the Meta page to reach the attacker's site. That site was hosted on sites.google.com, Google's legitimate site-building platform, which adds a second layer of abused infrastructure.

The Google-hosted site was most likely built to harvest credentials. However, when At-Bay Security's team tried to access it to confirm what the attacker was after, Google had already taken it down. That suggests Google is actively finding and removing abusive pages on its platform.

Why this works: Because the attacker used the actual platform of a legitimate, trusted company to deliver or host a message, the email can successfully get through email security filters. This tactic works by avoiding the specific signals email security tools rely on, since a link hosted on a legitimate platform inherits that platform's own trusted reputation. If the targeted employee inputs their credentials in the attacker-created page, the credentials go straight to the attacker instead of Meta, handing over control of the business account, including its connected ad accounts, pages, and any payment methods on file.

How to prevent it:

- **Train employees to verify identity, not just infrastructure.** The email passed authentication checks because it really was sent through Meta's system, but genuine infrastructure doesn't guarantee a genuine request. Having employees confirm an unexpected business partner request through a known contact or channel before accepting it closes off this exact laundering trick.
- **Train employees to judge the page by what it asks, not where it's hosted.** A page on sites.google.com carries Google's reputation, not the actual publisher's. If a site built on a trusted platform suddenly asks for a login to a different service (Meta credentials, in this case), that mismatch is the signal, even when the domain itself is clean.
- **Require phishing-resistant MFA on business platform accounts, not just email.** These attacks target the credentials for advertising, business, and social accounts, which often carry weaker protections than a corporate email login. A passkey or FIDO2 key closes off this style of attack the same way it does for email-based phishing.
- **Use detection that reads context, not just reputation.** Domain and sender-reputation checks miss attacks that ride on infrastructure everyone already trusts. Behavioral and contextual analysis, which flags what a message or page is actually doing rather than where it came from, is what catches this kind of attack after the reputation check passes.
- **Treat unexpected outreach as a signal on its own.** Most of these attempts start and end with a single email — nothing the employee asked for, signed up for, or was expecting beyond that one unsolicited message. An unprompted request to log in or view a document is worth pausing on before anything else about the email gets a second look.

At-Bay Stance MDR for Email



At-Bay's MDR for Email watches for the brand impersonation, abused infrastructure, and credential-phishing patterns described above, flagging them before an employee has a chance to click. Combined with 24/7 monitoring from At-Bay Security, it's designed to catch the messages that get past a standard filter.

Learn more about MDR for Email →

CHAPTER 2

Remote Access and VPN Rankings

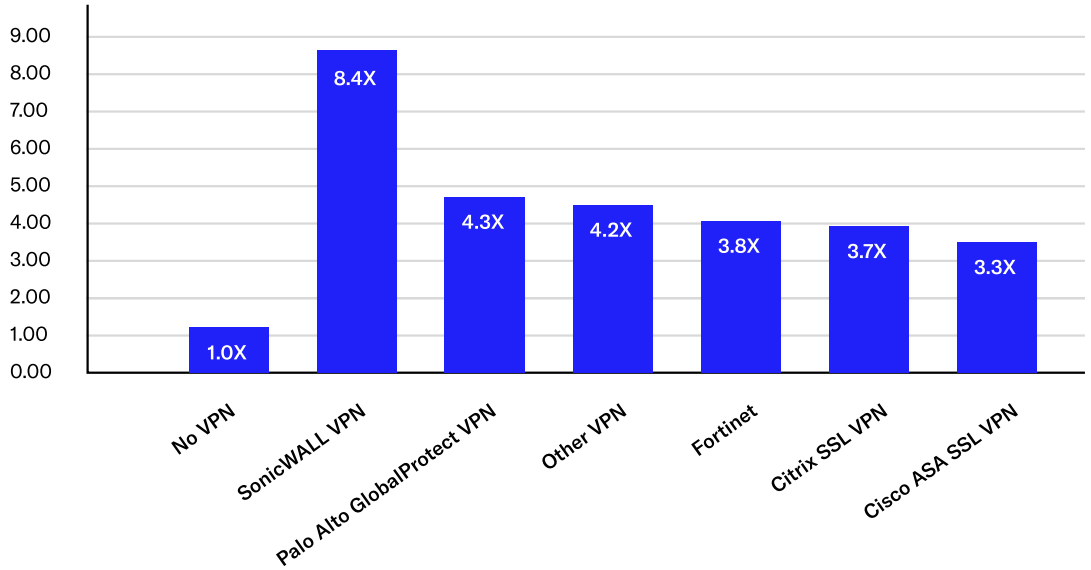
As we saw in the last chapter, email attacks aren't slowing down and credential harvesting is a major focus for threat actors. But what happens next? The real damage starts once that harvested credential gets tested against an exposed VPN portal and the attacker gains access to the entire network.

VPNs continue to be correlated with significant risk. Year after year, our claims data has shown that businesses with any kind of VPN detected fare worse than businesses without one, a pattern that holds again this year. All vendors in our current rankings, even the ones that improved, still show meaningfully higher risk than the no-VPN baseline.

FIGURE 9

Likelihood of a Business Using One of These VPNs to Fall Victim to a Ransomware Attack

Indexed Ransomware Claims Frequency by On-Premise VPN, 2025 - Q1 2026



65% of At-Bay Security ransomware investigations with a known initial access vector relate back to VPN and RDP. Of those, 21% of cases are associated with a SonicWall VPN.

Customers using SonicWall saw a significant increase in risk year-over-year, up from 5.8X to 8.4X. We flagged this trend building in last year's report: In Q3 of 2025, At-Bay's Response & Recovery team saw a 300% jump in Akira ransomware cases quarter over quarter with the vast majority of them using SonicWall devices as the observed entry method.

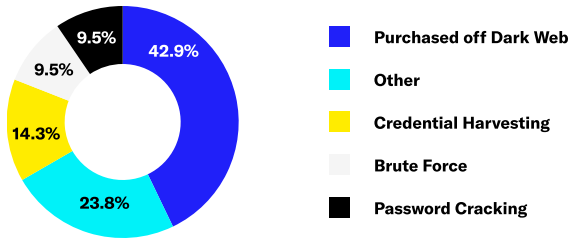
While many details about Akira's maneuvers remain unknown, based on our own At-Bay Security data, we see that most initial network access started with a login through a SonicWall SSL VPN using legitimate credentials, not a novel exploit. We cannot always confirm how threat actors obtain credentials, but it's clear that they're highly sought after by threat actors.

In extortion incidents where At-Bay policyholders paid a ransom and received a security report from the threat actor, the threat actor frequently cited credentials or techniques used to identify credentials as the root cause for the attack. In nearly 60% of cases, credentials were allegedly harvested by the attacker or purchased off the dark web.

FIGURE 10

Threat Actors Cite Credential Theft as Most Common Entry Vector

Initial Access Methods Reported by Threat Actors to At-Bay Security



"Credential Harvesting" refers to instances where threat actors stole the abused credentials themselves, while "Purchased off Dark Web" means that some other individual or group stole the credentials, posted them on the dark web, and the threat actor who launched the attack bought them in order to abuse them.

This is the throughline for this year's report: credential theft is not just an email security problem, and remote access exposure is not just a patching problem. They're often two ends of the same attack chain. A business can have a VPN and still be breached because an employee's credentials were phished. Conversely, a business can have excellent email security and still be breached if a leaked credential from an unrelated source is tested against an internet-facing login portal with no MFA behind it.

8.4

How Credential Theft Leads to Stolen Data

1

EXPOSED

Credentials leaked online

At-Bay Security observed a case in which login credentials for multiple users were posted in plain view on a Telegram channel.

2

DORMANT

Three weeks go by, unused

The exposed credentials sat untouched for about three weeks before an attacker put them to use.

3

ACCESSED

Threat actor logs into AWS

Using several of the compromised accounts, the attacker logged into the policyholder's public-facing Amazon Web Services (AWS) portal and ran database queries.

4

EXTORTED

Data stolen, then leaked

The queries surfaced records with data, which the attacker stole to extort the company. When the extortion demand went unmet, the attacker posted the database's contents to an underground forum.

CASE STUDY

At-Bay Stance MDR for Identity Blocks 1 a.m. Akira Attack

The Threat

Akira launched a VPN exploitation attack at 1 a.m. on a Saturday morning, when most IT teams are offline. The attacker moved fast, going from initial access to lateral movement in less than 5 minutes. The attacker's lateral movement was their downfall: At-Bay MDR identity tools were live and monitoring for anomalous logons.

The Response

At-Bay's MDR team was watching. Analysts spotted the intrusion, quarantined the affected systems, and ejected the attacker in under 10 minutes. They then removed the attacker's tools and fixed the compromised credentials. Last, they ran a full forensic review of file writes, registry changes, and event logs to confirm the environment was clean.

The Result

Akira never finished reconnaissance or staged ransomware. Once notified, the customer restricted access to the compromised VPN, reset accounts, and patched, ending the day more secure than it started.

Why it Matters

Stopping an attack at 1 a.m. takes a team that's watching the entire environment 24/7 and can take immediate hands-on action to shut it down.

Learn more about MDR from At-Bay at at-bay.com/mdr

Response timelines differ. Past results do not guarantee future outcomes. This content is provided for information purposes only and is not intended to define any Policy commitment. No warranty is given or liability accepted regarding this information.

At-Bay Stance MDR is provided by At-Bay Security, LLC, and available to eligible businesses with or without an insurance policy placed through At-Bay Insurance Services, LLC. At-Bay Security, LLC is a wholly owned subsidiary of At-Bay, Inc., providing cybersecurity services including MDR and incident response.

FortiBleed

On June 16, 2026, cybersecurity vendor SOCRadar announced it had discovered a trove of more than 80,000 sets of working credentials for Fortinet devices[23]. The credentials appeared to match devices located in 194 countries, with many credentials verified as authentic at the time of the discovery. SOCRadar described the credentials as the accumulated loot from a deliberate and long-running credential harvesting campaign targeting Fortinet network security appliances, with SOCRadar attributing responsibility for the campaign to the Lynx / INC ransomware group[24].

Anticipating that the FortiBleed discovery heralded a period of increased risk for Fortinet customers, we worked to quickly inform our policyholders of the risk and the appropriate response necessary to protect their environment from compromise due to stolen credentials.

At-Bay policyholders with Fortinet devices appear to have avoided the worst of the fallout from FortiBleed, but there are many aspects of a large credential leak like FortiBleed that subject the enterprise to risks that are much longer term and more pervasive than might be implied by a typical compromised credentials scenario:

1. **Security researchers have not definitively identified the source of the credentials in the FortiBleed database.** This means that we can't verify that the vulnerability or set of vulnerabilities that enabled the credentials to be stolen in the first place have been remediated, so the credential harvesting may be ongoing among Fortinet customers.

2. **The database containing the stolen credentials may not include all credentials that have been compromised as part of the campaign.** The trove discovered by researchers may have been a subset of the total. Many At-Bay policyholders were relieved to find that their credentials were not in the database discovered by SOCRadar, but that should not be taken as proof that they weren't compromised.

3. **Compromised credentials for Fortinet devices may open doors elsewhere, including on Microsoft Active Directory instances or on unrelated SaaS platforms.** At-Bay advised policyholders who identified their credentials in the FortiBleed campaign to reset all credentials in their environment due to the risk of reused credentials across applications and services. For companies that did not do this, credentials from Fortinet devices may therefore remain useful to attackers long after victims have reset credentials on their Fortinet devices or even removed them from the network altogether.

Companies operating Fortinet appliances on their network are almost 4X as likely to fall victim to a ransomware attack compared to those companies with no identifiable remote access infrastructure. The risk of these appliances was already high before they were implicated in a massive credential harvesting campaign, which could result in the compromise of almost any other application, service, or system that a company operates. While Fortinet has offered configuration recommendations to operate their appliances safely, many companies simply aren't capable of following them. At-Bay continues to recommend that organizations operating on-premise network appliances that enable remote access migrate to more modern and lower-risk cloud-based remote access tools such as zero-trust solutions.

[23]Source: SOCRadar, FortiBleed: SOCRadar's Investigation into 86,644 Compromised Fortinet Firewalls, June 2026
[24]Ibid.

Why a Stolen Credential Can Hand Over the Whole Network

A stolen credential should be a narrow problem: one password, one account, caught early in a well-defended environment. However, on most of the VPN appliances in our claims data, it's the whole network instead. Cisco, Citrix, SonicWall, Fortinet, and similar VPN products all do the same job: authenticate a credential, grant network access. Without enforced MFA, conditional access, or anomaly detection on login patterns, whoever holds the credential holds the network.

The Akira ransomware cases observed by At-Bay in 2025 show how little room that leaves for defenders: average dwell time from login to encryption was 1.5 days, and in one case, the entire intrusion occurred in under 40 minutes.

Endpoint detection and response (EDR) was often in place but was rarely enough on its own. Most At-Bay policyholders with EDR deployed were still fully or partially encrypted, except where EDR was actively monitored by a team ready to act within minutes. That distinction, not just the tool, is what contained the intrusion. (Notably, no At-Bay MDR customers were encrypted during Akira's campaign.)

Credential theft itself is hard to stop outright. What's fixable is making sure a stolen credential isn't enough on its own: enforced MFA, actively monitored detection, and, where possible, a move away from architectures that grant network-wide access on a single login.

Real World Findings From At-Bay PCEH

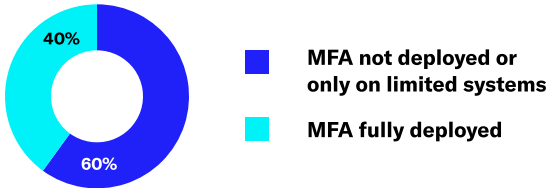
In April 2026, At-Bay introduced a new security service offering to our policyholders who have experienced claims — Post-Cyber Event Hardening (PCEH) — which offers support with hardening the technology environment against future cyber attacks. Picking up where incident response and recovery activities conclude, PCEH aims to rapidly identify the likely causes of the next cyber incident and remediate them with the intent of improving the overall security posture of the customer.

In delivering PCEH services over the past six months, the At-Bay Security team has surfaced insights about the cyber posture of companies that have experienced damaging cyber attacks. Specifically, the PCEH team has identified the top control deficiencies likely to be present among companies that have experienced a cyber event even after they have completed their necessary incident response and recovery.

FIGURE 11

60% of Policyholders Had MFA Missing or Insufficiently Configured / Deployed

According to At-Bay PCEH Data, 2026

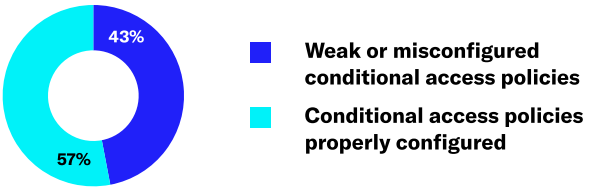


[26]Figure 9: Likelihood of a Business Using One of These VPNs to Fall Victim to an Attack
[27]Source: Fortinet, Analysis of Reported Credential Compromise of FortiGate Devices, June 2026

FIGURE 12

57% of Policyholders Had Weak or Misconfigured Conditional Access Policies

According to At-Bay PCEH Data, 2026



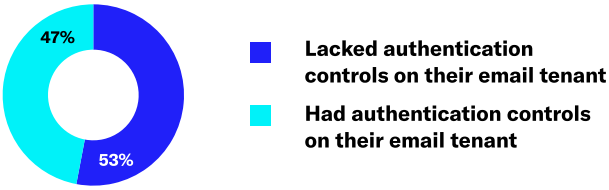
57% of policyholders that used At-Bay's PCEH service had weak or misconfigured conditional access policies. These policies, deployed on network devices such as firewalls and remote access infrastructure, typically enforce access controls by denying connectivity to users attempting to access critical systems, applications, or network segments unless they meet certain criteria related to the user's role, assigned privileges, or the configuration of the user's device.

The most common conditional access policy gap was the lack of enforcement of MFA usage among users who were already enrolled in MFA. Essentially, these users could opt to authenticate without the company's established MFA solution, nullifying the expected impact of deploying MFA in the first place. The remediation for these policyholders included enforcing MFA usage for all users and ensuring that all access points to business critical systems and data were operating with the same access control policies.

FIGURE 13

53% of Policyholders Lacked Authentication Controls on Their Email Tenant

According to At-Bay PCEH Data, 2026



53% of policyholders that used At-Bay's PCEH service lacked authentication controls on their email tenant. These controls, which include Sender Policy Framework (SPF), DomainKeys Identified Mail (DKIM), and Domain-based Message Authentication Reporting and Conformance (DMARC), ensure that messages from a domain can't be spoofed, that messages to a domain can have their source authenticated by the recipient, and that a domain cannot be used for fraud and email scams. Not having these controls in place leaves a company vulnerable to inbound spoofed messages and may increase the risk of fraud. Organizations might also see their emails to third parties spoofed and might therefore be implicated in fraud affecting other companies. The remediation for these policyholders included configuring email authentication controls.

Remarkably, while these controls have a significant impact on the likelihood of a future cyber incident, implementing them typically carries no incremental financial costs for an organization.

The reason for this is that they're already available in the security tools that many organizations have currently deployed. The additional controls just need to be activated and configured properly to become effective. At-Bay will continue to share insights from our PCEH team as the service gains further uptake among our policyholders who experience claims.

Conclusion

In this report, we illustrate how a single weakness can be used to exploit an entire system. A password captured on a spoofed login page sent via email doesn't just expose one inbox. It can be tested against a remote access portal, reused on a business platform, or sold to someone else who will try it everywhere it might work. The examples in this report show how an attack that starts in one system can find its way into another, and how a gap that looks small on its own can become the first step toward something much larger.

That means understanding the effectiveness of the tools ranked in this report continues to be a top priority for us. The tools a business chooses for email and remote access may still be a determining factor in the amount of risk a business takes on.

Choosing the right tools — and ensuring best security practices are being applied — for each part of the environment matters.

Understanding and managing the complexity of today's cyber risk takes a holistic approach to defense, and that's where InsurSec comes in. Most security vendors see attacks from the perspective of the tools they sell. As an InsurSec provider, At-Bay sees what happens when attacks succeed against our insured businesses: which tools were in place, how attackers got in, and what it cost. We turn those claims insights into threat intelligence, and that intelligence shapes our security products, from the signals our MDR team watches for to the alerts we send when a new threat emerges.

Insurance and security each make the other stronger, which is why we believe InsurSec is the most comprehensive approach to cyber risk available today.

Methodology

At-Bay's analysis is based on claims information collected from insureds with policies placed through At-Bay Insurance Services, LLC (“At-Bay Policyholders”) from 2021 through the first quarter of 2026, with the email and remote access rankings in this report reflecting claims from 2025 through Q1 2026. Claims reviewed included incidents related to malicious emails, financial fraud, remote access, and ransomware. By analyzing actual claims data alongside data from At-Bay Security's investigations and services, the At-Bay Research team set out to answer these questions:

- How are email attacks and email defenses changing?
- How do outcomes associated with specific email solutions and email security solutions differ?
- How do remote access solutions change the risk profile of a company?
- How are email and remote access risk connected through credential theft and abuse?

Email Analysis and Rankings

To evaluate email risk, our researchers established an indexed email claims frequency for each solution and compared it against the average for all At-Bay customers using that category of solutions. To show the longer-term trend, overall email claims frequency was also indexed to the average overall frequency in 2021. Solutions were analyzed in two categories. The "Email Solutions" category covers Google and Microsoft 365, the two most prominent email providers tracked by At-Bay. They aren't email security companies, but the security features built into their platforms set a baseline of protection, and for many organizations without a separate email security product, that built-in security is the only email security they have. For that reason, we track their performance side-by-side with email security vendors. The "Email Security Solutions" category covers six secure email gateway providers prominent among At-Bay policyholders: Proofpoint, Mimecast, Barracuda, AppRiver, Intermedia, and Sophos.

By identifying the solutions with high or low claims frequency compared to the average, we believe we can assess the relative effectiveness of each solution at reducing the risk of incidents that start with email. To help explain changes in the rankings, we also reviewed the security capabilities each vendor released during and just after the reporting period. Capabilities released late in the reporting period or after it ended (including several 2026 updates from Mimecast, Barracuda, and Sophos) may not be reflected in this year's results, and we expect to see their impact in next year's data. At-Bay doesn't typically publicize or rely on vendor-reported data, and the rankings in this report are based on our own claims data.

Brand Impersonation and Abused Infrastructure Analysis

The analysis of brand impersonation and abused infrastructure is based on more than 2.2 million messages flagged by At-Bay's MDR for Email between February 1 and August 31, 2026. Flagged messages were categorized by tactic (brand impersonation or abused infrastructure), by the brand or service involved, and by the attacker's primary intent. Messages without a classified intent were excluded from the intent analysis.

Remote Access Analysis and Rankings

Our data has shown again this year that businesses with any kind of VPN detected fare worse than businesses without one. For this analysis, we compared the ransomware claims frequency of businesses using specific on-premise VPN solutions against a baseline of businesses with no VPN, indexed to 1.0. Solutions analyzed include Cisco ASA SSL VPN, Citrix SSL VPN, SonicWall, Palo Alto GlobalProtect, Fortinet, and an "Other VPN" group.

To be clear, our claims data does not show that these products are directly responsible for every claim. In ransomware investigations conducted by At-Bay Security where the initial access vector was known, VPN and RDP accounted for 65% of cases, which means some incidents began elsewhere. In many of the cases this report highlights, attackers entered through a VPN by logging in with legitimate, stolen credentials rather than by exploiting a vulnerability in the product itself.

A Note About Our Additional Data Sources

Beyond At-Bay Claims data, this report draws on At-Bay Security incident investigations, security reports that threat actors gave policyholders who paid a ransom, and findings from At-Bay's Post-Cyber Event Hardening (PCEH) service during its first six months after launching in April 2026. PCEH findings reflect only policyholders who used the service after a claim. The figure showing that 30% of policyholders believed they had MFA activated when they didn't comes from a separate analysis from At-Bay's Risk team of claims over the past four years.

Contributors



Adam Tyra
CISO for Customers



Cameron Pitts
Senior Content Marketing Manager



Laurie Iacono
Director, Threat Intelligence



Michael Lowe
VP, Head of Marketing



Mike Scutt
VP, MDR Security Operations



Samantha Wong
Senior Risk Analyst

About At-Bay + InsurSec Report

At-Bay is the InsurSec provider for the digital age, helping businesses mitigate cyber risk and avoid incidents by continuously analyzing data from security scans and collecting cyber threat intelligence and the relevant details of security incidents reported by insureds. Because we can correlate information about a significant number of real-world incidents with data about the victim's technology environment before the incident occurred, this enables us to reliably identify trends and relationships that other companies and security vendors cannot. We're able to clearly identify security controls that mitigate risk, differentiate them from security controls that don't mitigate risk, and prove our case with empirical data from actual incidents where those security controls were in place.

Our goal is to share our findings on the respective impacts of a range of security controls with the public at large. We believe we can use facts and evidence to cut through the noise of a crowded cybersecurity marketplace and enable organizations to deploy scarce cybersecurity resources for maximum impact. We regularly develop and share a slate of statistically provable leading practices for security that can be readily consumed by organizations regardless of headcount or the size of their security budget.

The information contained is for general guidance on matters of interest only and is not intended to construe or the rendering of professional services of any kind. If professional advice is required, the services of a professional should be sought. All information is provided as is with no guarantee or warranty of any kind, express or implied, concerning the completeness, accuracy, usefulness, timeliness of the information provided. At-Bay is not responsible for any errors or omissions, or for the results obtained from the use of the information provided in these materials. This report post includes links to third-party websites. These links are provided as a convenience only. At-Bay does not endorse, have control over, or assume responsibility or liability for the content, privacy policy, or practices of any such third-party websites. At-Bay Insurance Services LLC, a wholly owned subsidiary of At-Bay, Inc., is a licensed insurance agency and surplus lines broker in all fifty states and the District of Columbia.

©10/2026 At-Bay. All Rights Reserved.