TRAVELERS

Q1 2026

Cyber Threat Report

The Growing AI Governance Gap

Table of Contents

Published by Travelers with contributions from:

Ryan Bell
Director, Threat Intelligence – Cyber Risk Services
Travelers

Rehman Khan
AVP, Cyber Risk Services
Travelers

Christine Mapes
Managing Director and Counsel, Claim
Travelers

Courtney Bergeron
Cybersecurity Technologist
Travelers

Matthew Butler
Director, Cyber Risk Services
Travelers

Alex Pinto
Sr. Director of Product Marketing – Cyber
Travelers

Executive Summary

When last year ended with the highest level of ransomware activity we've observed, the natural question was whether that surge would prove durable. The answer, unfortunately, appears to be yes: the first quarter of 2026 recorded 2,405 victims posted to leak sites, essentially flat quarter over quarter, and the second-highest quarter we have tracked.

The story behind the numbers is one of fragmentation. Eighty-four distinct ransomware groups were active in Q1, which is the highest count in our dataset going back to 2020 — and 19 of them made their first appearances on leak sites during the quarter. Qilin remained the most active operator for the second quarter in a row, but newer entrants like the "Gentlemen" group, which posted 207 victims just months after first appearing, illustrate how quickly the landscape is reshuffling. This presents a challenge for organizations hoping to get a grip on the prevailing tactics and approaches (and how to defend against them).

This quarter we also turn our attention to the proverbial elephant in the room. AI is starting to impact overall risk for many organizations, though perhaps not in the ways you might expect. While recent coverage of Anthropic's Mythos model has reignited speculation about AI-powered attack chains, the more immediate exposure for most organizations is not what threat actors are doing with AI, but what their own employees are doing with it. We offer some guidance to organizations seeking to develop a governance approach for AI use.

Ransomware levels remain elevated: 2,405 victims were posted to leak sites in Q1 2026, the second-highest quarter we've tracked and just 2% below Q4 2025's all-time high.

The ecosystem fragments further: 84 distinct ransomware groups were active in Q1 — the most in our dataset — with 19 making their first leak-site appearance. New entrant "Gentlemen" ranked second in overall volume.

The most actionable AI risk is internal: At least 43% of U.S. workers now use AI on the job, but governance models built for conventional software don't capture how AI tools are used or what they are capable of. Developing governance is a priority for any organization.

Ransomware Activity Remains Elevated

The first quarter of 2026 recorded 2,405 victims posted to ransomware leak sites,* maintaining activity at near-record levels for the second consecutive quarter. While this represents a 2% decline from Q4 2025's all-time high of 2,453 victims, the more significant story is what this figure tells us about the broader trajectory. Historically, activity has dropped sharply after a peak quarter. This time, the heightened level of activity stayed put.

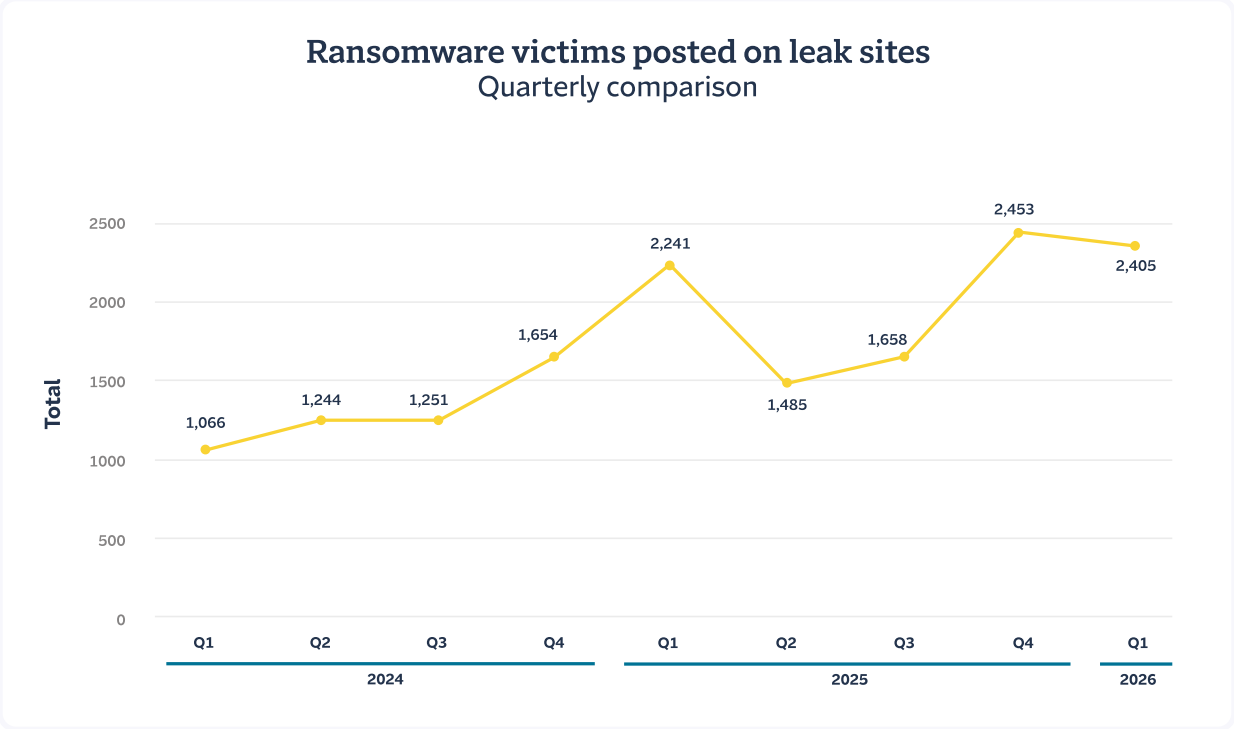

Year-over-year, Q1 2026 was up 7% compared to Q1 2025's 2,241 victims. That prior quarter was itself a localized peak — it was the highest Q1 we had recorded at that time — so continued growth above that level reinforces what has become a consistent theme in these reports: ransomware activity is increasing, with each new high watermark tending to set the floor for what follows.

*Data shared on leak sites provides a proxy for overall ransomware activity. A victim's information will typically be posted if the victim has refused to pay a ransom. This means that the data should be viewed as a fraction of overall activity, but one that can provide a longitudinal comparison of activity over longer time frames.

The word of the quarter: "fragmentation"

One of the notable features of Q1 2026 was not just the volume of attacks, but the number of groups carrying them out. Eighty-four distinct threat actor groups were active during the quarter — the highest count yet in our dataset that goes back to 2020. That compares to 70 in Q1 2025 and 63 in Q4 2024. It's also notable that 19 of the groups active in Q1 made their first appearance in leak site data that quarter.

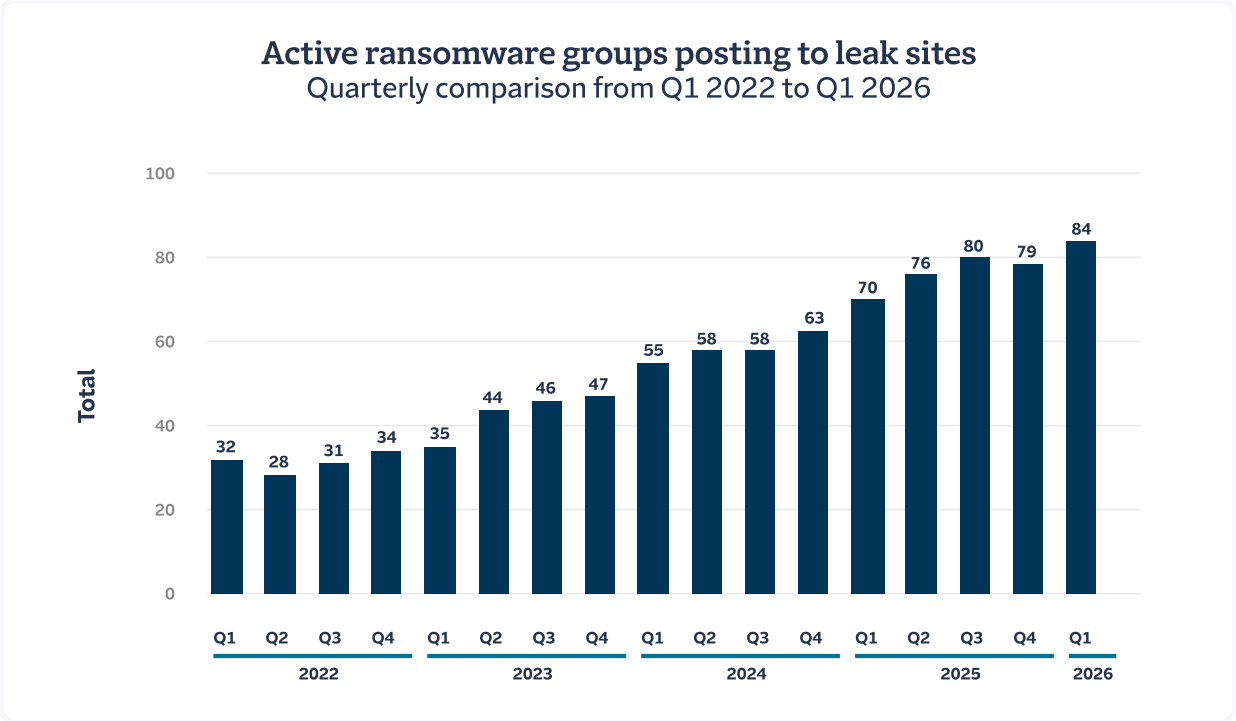

We've written in previous reports about the tendency of ransomware groups to cycle in and out of activity. What's different now is that the churn is accelerating on both ends: new groups are entering the ecosystem at a faster pace, while established groups are also dropping off at a higher rate. In Q1 2026, 20 groups went inactive, while 19 new ones simultaneously made their debut. The net effect was still a record number of active groups — but it also points to an ecosystem that appears to be growing more competitive and fragmented, not less.

From a risk perspective, a more fragmented ecosystem is not necessarily a less dangerous one. In some respects, it is harder to manage. Disrupting a single dominant group — as law enforcement has done in past operations against LockBit, RansomHub and others — becomes less impactful when attacks are distributed across a larger number of operators. Even as individual groups rise and fall, the total volume of attack activity shows little sign of abating.

The most active groups hold their share

As you'd expect given the increase in active groups, the share of overall attack activity by the most active groups declined – although only slightly. In the first quarter of last year the three most active groups collectively accounted for 37% of all leak site postings, while in Q1 2026 that figure dropped to 34%. Qilin remained the most active group for the second consecutive quarter, posting 414 victims in Q1 2026. While this was a pullback from its Q4 2025 peak of 551, it still represented a 273% increase over its Q1 2025 activity. Akira, INC Ransom, PLAY and DragonForce continued as consistent contributors to the overall volume. (For a closer look at Qilin's history and tactics, see our Q4 2025 report.)

New groups appear

Among the 20 new groups making their first appearance in Q1 2026, one did so with a level of activity that would be notable even for an established operator. The group operating under the name "Gentlemen" posted 207 victims in the quarter — the second-highest total of any group — despite having only appeared in leak site data for the first time in September 2025.

In the span of roughly six months, Gentlemen went from no observable presence to the second-most prolific ransomware group in the world by this measure. That trajectory is unusual. Most groups that reach high volumes do so over multiple quarters, building infrastructure, recruiting affiliates and refining their tactics. Gentlemen's rapid ascent may reflect a group with significant prior experience operating under a different name, a well-funded new entrant that was able to acquire or develop capabilities quickly, or both.

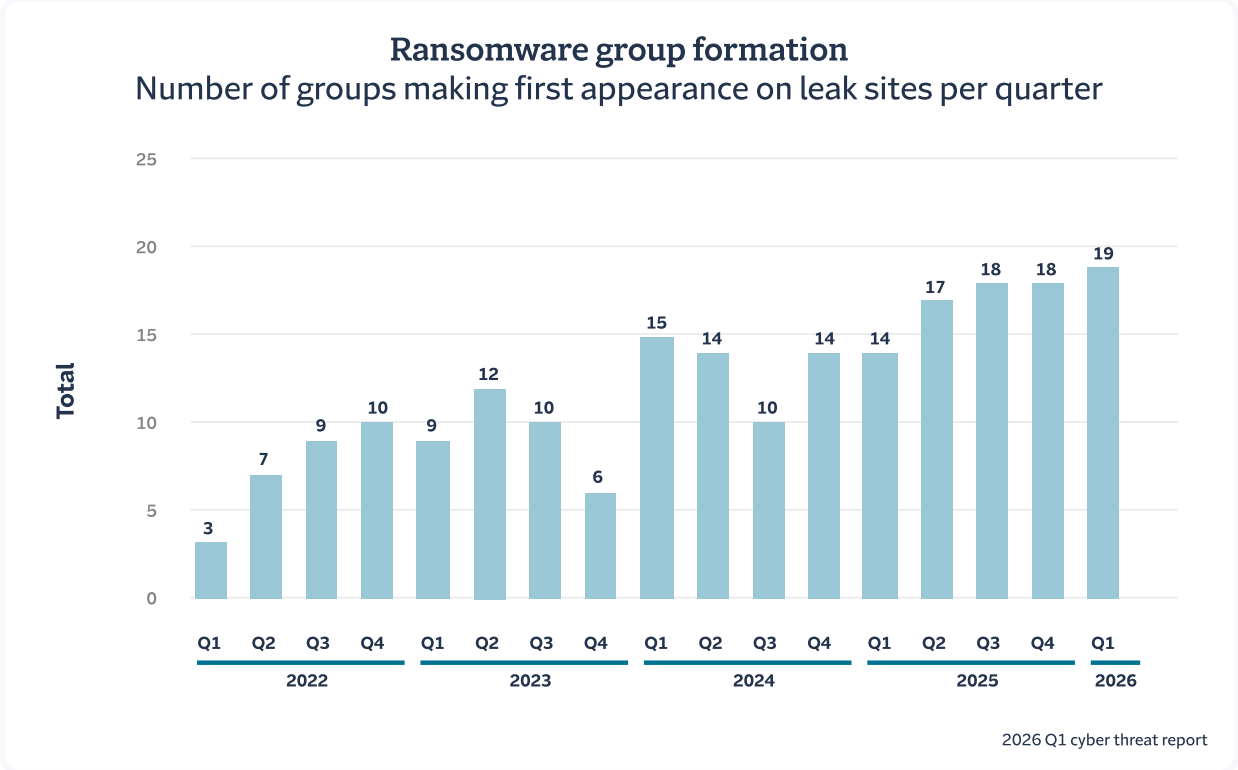

AI impacts ransomware – but not always how you expect

From practically the moment ChatGPT was released in November 2022 there has been fear over the application of large language models (LLMs) to ransomware. (This report from Check Point Research was out within three weeks). Visions of automated attacks abounded.

But since then, the real impact of AI usage by cyber criminals hasn't been easy to disentangle from evolutions in the cybercrime ecosystem that likely would have occurred with or without the advent of LLMs. The main drivers behind changes in ransomware activity we have pointed to in past reports – new software vulnerabilities, geopolitical disruption, law enforcement activity, the development of the RaaS ecosystem – have not centered on AI.

Then artificial intelligence company Anthropic announced a model that it says is better than almost all human researchers at discovering software vulnerabilities, and can do so much faster than them. It has not released the model to the public out of concern that it could be exploited by bad actors. Despite the careful approach to the release, the news touched off another cycle of speculation. it's worth game-planning possible scenarios as AI capabilities evolve, and the news certainly provides a great reason to make sure organizations are up to date on recommended cyber hygiene. But the actual evidence of AI use by adversaries so far points to less obvious impacts than the fully automated attack chains that keep cybersecurity professionals up at night.

One of these emerging areas of AI impact is at the negotiation stage.

(continued on next page.)

The group's targeting profile is also worth noting. While most of the top ransomware groups are heavily concentrated on U.S.-based victims, Gentlemen's activity was more globally distributed. Its most targeted countries in Q1 2026 included Thailand, the United States, France, Brazil and Turkey. That geographic breadth, combined with its targeting of financial services firms, hospitals, government agencies and IT providers, suggests a less discriminating approach that maximizes the pool of potential victims.

This lack of discretion illustrates a broader theme. In recent quarters, the barrier to entry into the ransomware ecosystem has continued to fall. The proliferation of ransomware-as-a-service (RaaS) platforms means that groups with the necessary criminal connections can acquire sophisticated tools and operational infrastructure without having to build them from scratch. The rapid displacement of established operators by newer entrants — and the simultaneous presence of 84 active groups in a single quarter — are symptoms of the same underlying competitive dynamic.

A threat landscape measured in dozens of active groups, rather than a handful, is more difficult to anticipate and prepare for. Monitoring the tactics of the largest groups remains important, but it provides an incomplete picture when a significant share of attacks is coming from actors that weren't on anyone's radar just a few months ago. And this variety layers onto the simple fact that overall ransomware activity has tripled since 2022 based on leak site activity, and ransomware claims have increased 80% since that year according to Travelers' internal claim data.

How Rapid Adoption of AI is Affecting Risk Exposure

AI impacts ransomware – but not always how you expect: (continuted)

Negotiation specialists we work with have reported seeing chatbots take the place of human interlocutors in negotiating a ransom. The evidence that they are not human comes partly from conversational style, but more definitively from the speed of responses – particularly when actions, such as the removal of data from a leak site, occur instantaneously upon prompting.

These bots may be "agents" and therefore part of the recent trend of agentic AI. (These days the definition of what constitutes an agent is hotly debated.)

In any case, they are clearly the result of improvements in AI models that allow for more complex, multi-stage instructions to be reliably followed.

This leads to some interesting implications. Negotiators have always dealt with things like language barriers, inconsistencies in negotiation strategy and even time zone incompatibility (cyber criminals sleep, too). Human foibles have a tendency to crop up even in the most impersonal and adversarial exchanges. A bot instructed not to accept any number below a certain threshold, however, will not waver.

If bots programmed with strict instructions are taking over negotiations, how does that change a negotiator's game plan? It may remove an opportunity to play on the emotions – or energy level – of the attacker, but it may also introduce other tactical possibilities. Bots may prove to be fallible in different ways than humans.

We'll need to see more evidence and data to say how this impacts the end results of negotiations, but it's something we are tracking closely.

Discussion of Anthropic's Mythos model has dominated the topic space around cybersecurity and AI recently, because of the potential that it (or more likely its expected imitators) might have on threat actors' attack capabilities. We are watching closely for attack activity that displays the use of advanced AI models; if we see a change in common attack patterns, you can be sure we'll report our findings. But for now, if you're not a software company there is little to do in the realm of perimeter defense that's different than what we would have recommended a year ago – or even a few years ago. (Things like having a vulnerability management program to help efficiently apply software patches, for instance.)

Instead, when we're asked by policyholders and their brokers, "how should we think about AI risk?" our answers focus on what organizations do have control over right now. Mostly that means **how they are managing the deployment of AI in their organizations.**

Rolling downhill on rollouts

In August last year the Federal Reserve Bank of St. Louis found that 37% of Americans were using AI for work. The bank reviewed similar surveys from the past and found that at the same point in the tech lifecycle – roughly three years after a major commercial release – adoption of personal computers for work stood at 25%, and adoption of the Internet for any purpose (personal or work) was 30%.

The upshot: **adoption of AI is happening more rapidly than any previous workplace technology and does not seem to be slowing.** By early 2026, when the bank ran another survey, workplace AI usage had already risen to 43%.

In conversing with policyholders, we're seeing a wide range of approaches behind the scenes of those striking adoption rates. Some are resisting the tidal shift while others act as laboratories for a litany of AI experiments. All the approaches are well-intentioned in their own ways, but also carry their own downstream consequences.

The following are three rough groupings of the approaches we're seeing.

Approaches to governing AI deployments vary widely

The Conservative Approach

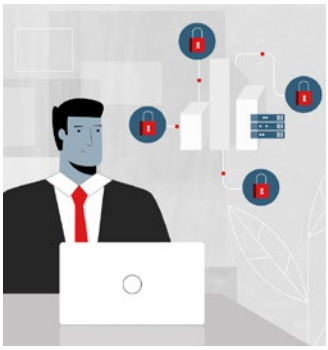

Official position	First order consequence	Second order consequence
No official adoption of AI company-wide, or limited pilot to small number of employees.	Ungoverned shadow usage by employees is likely. Even if restrictive web filtering is in place on the company network to prevent usage, employees may be exporting data to personal devices to use AI.	No visibility into potential data leakage since employees try to conceal their non-compliance (or it occurs off network).

The Enabler Approach

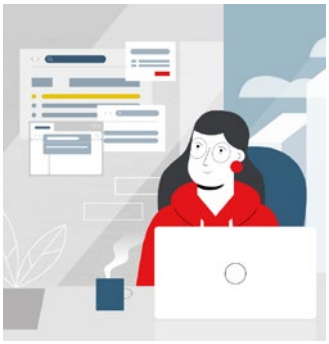

Official position	First order consequence	Second order consequence
Allow employees to use whatever tools they see fit, to whatever extent they see fit.	A dizzying number of overlapping tools are in use, each with access to different sets of data.	Unpredictable interactions between tools leads to unintended outputs or exposures.

The Early Adopter Approach

Official position	First order consequence	Second order consequence
Sanctioned, controlled rollout of specific tools – but with few limitations on use of the sanctioned tools.	Some employees begin experimenting with leading-edge capabilities, like AI agents.	Effects that weren't anticipated – such as AI agents doing things their creators didn't intend, or simply losing track of the multiplying number of agents.

Depending on their approach to AI utilization, organizations will want to take steps to understand and manage risk. But getting started can be challenging. Unlike conventional software, AI tools are open-ended and probabilistic (rather than deterministic), so they don't map cleanly onto governance frameworks that were made for software that behaves the same way for all users all of the time. When it comes to governing AI, a lot depends on the nature of the business and what the employees do. It has to be tailored.

That is why the best approach is to start from the ground up rather than attempting to copy and paste an existing process or framework. The following are a few of the areas we recommend looking at first.

Build the governance foundation

If no one owns it, no one will manage it. Without defined roles and responsibilities for AI in the organization or a named individual or committee with accountability for decisions, any efforts to create or enforce governance will flounder. This is the place to start.

With the people in place to do the governing, they can turn to the policy side.

Documented policies and acceptable use guidelines should outline how employees can use AI tools responsibly across several dimensions, such as data handling, output validation and intellectual property considerations.

Compared to a tool like email, the spectrum of familiarity and trust with AI within an organization can be enormously broad. Some bold experimenters might be allowing AI to access every documentable bit of information about themselves in an attempt to automate their entire job. Others might have heard stories about hallucinations and won't touch AI tools out of principle when it comes to their work.

When thinking about an organizational policy, business leaders should think about meeting each of these varying perspectives where they are. Forceful, restrictive language won't help getting laggards on board (or might scare others into off-network shadow usage). Vague, permissive language will provide openings for the most experimental to push the envelope beyond what was intended. Strive to provide clear, neutral language on what's permitted, what's prohibited and how to handle sensitive data within AI tools.

Drafting of these organizational policies should ultimately lead to an Acceptable Use Policy (AUP) and a requirement for employees to acknowledge it (at a minimum). These foundational steps should come before almost anything else is done with AI.

Recommendation Recap: Establish an enterprise AI governance program defining roles, responsibilities, and oversight mechanisms, including at least one individual or committee with accountability. This body should develop and publish policies and procedures, including an AI Acceptable Use Policy, that address data handling, output validation, intellectual property and escalation procedures.

Rolling out a policy

Leaders who manage a team of almost any size probably know this intuitively: publishing a policy – no matter how much they implore employees to read it – doesn't always change behavior. Training and reinforcement must follow. Some employees may balk at yet another training, but with technology that's this new and changing this fast, it's non-negotiable.

Leaders should start with covering the practical fundamentals: how to handle data appropriately when working with AI tools, how to validate outputs before acting on them, how to recognize when a result might be unreliable and what to do when something seems off. The goal isn't to make every employee an AI expert — it's to ensure that no one is making high-stakes decisions based on AI outputs they didn't know to interrogate appropriately.

That said, training must be responsive to the level of permission granted in your planned deployment. If a business is allowing software developers to build AI agents that can act autonomously, for example, a specialized training should cover how these agents behave in relation to the software development lifecycle. This is especially important if employees who are *not* software developers begin contributing applications they create with AI, since they won't have the same built-in instincts for what is risky about the code they're introducing to the system.

There are stories of AI agents gone rogue, doing things like deleting an entire production codebase. Preventing a disaster like that can and should be a motivating factor in training a workforce. But an example of risky behavior can be as mundane as an employee sharing a single internal document with an AI tool that isn't authorized to handle it. Organizations that haven't established a formal training requirement are relying on individual awareness of risks in these kinds of routine situations, which will be variable.

Recommendation Recap: Implement a mandatory AI training program for all employees who use or will use generative AI tools, covering topics such as data handling, prompt hygiene, output validation and recognizing AI limitations. Require periodic refresher training to keep pace with evolving AI capabilities and emerging risks

Keep a human in the loop

If one key to preventing rogue behavior is through training, the next is building a slowdown or stoppage into certain processes. If AI tools are involved in making high-stakes decisions, removing the human from that process creates exposure that's difficult to defend. AI outputs can be fluent, confident, convincing – and wrong. Without a qualified individual reviewing recommendations before they take effect, errors can move through a workflow undetected and with real consequences.

> If AI tools are involved in making high-stakes decisions, removing the human from that process creates exposure that's difficult to defend.

This isn't to say that an AI policy should dictate that AI *not* be used in high-stakes contexts. But it should clearly articulate where the handoff happens. Not every AI-assisted task carries the same risk profile, but every business has certain functions that carry elevated risk, and a mandatory review checkpoint is one of the more immediately actionable controls available. (It's also one of the harder gaps to explain away, to customers or regulators, after something goes wrong.)

Recommendation Recap: Implement a mandatory human-in-the-loop requirement for all AI-assisted decisions classified as high-risk or sensitive, ensuring a qualified human reviews and approves AI recommendations before they take effect.

Vet tools before deploying them

All of the advice above can apply to a company that is selecting a single, general purpose AI tool like ChatGPT or Claude to roll out in their organization. But we already know that in practice that is a rare situation.

The pace at which AI components are being injected into existing software tools is dizzying. And then there is the crop of new AI-based software tools that may be better suited than a general tool for certain tasks. The pace that employees wish to introduce these capabilities – or that the software companies are pushing them on their customers – can easily outrun regular evaluation processes.

Businesses should adapt their software procurement process to cover how AI tools handle sensitive data, the vendor's own security posture and what regulations could come into play. Hopefully these are all considerations that already go into making software decisions, but even if they are, the fact that AI tools have an element of unpredictability inherent in the way they operate changes the tenor of these decisions. It may not be as simple as defining access controls and restricting certain features.

The same logic applies to tools already in use. If a privacy risk assessment hasn't been completed on applications currently running in the environment, the organization may not have a clear picture of what data those tools are collecting, processing or sharing. That could include applications that have been in place so long that they pre-date the kind of risk assessment we're talking about – think email, PDF readers or spreadsheets. Embedding a privacy impact assessment into both the approval process for new tools and as a retroactive step for existing ones gives the organization a defensible baseline and a clearer view of where the risks sit.

Recommendation Recap: Integrate AI-specific risk criteria into the third-party risk management process, requiring completion before any AI tool is deployed. Conduct privacy impact assessments on all AI tools currently in use, and make privacy impact assessments a mandatory step in future AI tool approvals.

Conclusion

Once an organization's AI lead or committee gets into the flow of developing policies and procedures, the devil will be in the details. Some employees may feel like their AI usage is being constrained, others may feel like something is being forced on them. We wouldn't ask any organization to ignore its own employees. But efforts to accommodate every employee's preferences and use cases can sometimes backfire and lead to exposure. Hopefully this guide helps provide a starting point, but organizations should be sure to seek guidance from their legal or other professional advisors to determine their best approach to AI implementation and governance. For Travelers cyber policyholders, our Cyber Risk Services team is available to provide additional guidance and resources to help manage AI-related risks that could lead to exposure under your cyber policy.

Social Engineering Evolves to "Multi-Vector" Attacks

The basic vocabulary of social engineering has not changed much over the past decade. Phishing, spear phishing, vishing ("voice phishing" through a phone call) and business email compromise (BEC) have all been part of the threat actor's toolkit for years. (One exception is the introduction of deepfake technology, something we covered last year).

Looking at claim figures, it makes sense why things haven't changed much: these tactics remain damaging, even at a time when training on how to spot them has been embedded in organizations for years. According to Travelers internal claim data, the severity of claims that combine Social Engineering Fraud (SEF) and BEC is up more than 30% since 2023, and those two categories continue to make up roughly 40-50% of all cyber claims at Travelers.

According to Travelers claim data, the severity of claims that combine Social Engineering Fraud (SEF) and BEC is up more than 30% since 2023, and those two categories continue to make up roughly 40-50% of all cyber claims at Travelers.

What *has* changed about these attacks is the way tactics are being combined. We are now seeing more instances of coordinated sequences in which two, three or more tactics are stitched together. What does this "multi-vector" social engineering look like in practice? Below we walk through an illustrative pattern: the "mailbomb + ClickFix" combination. Members of our claim team selected this example because it reflects recent real-world situations that caught their attention.

Anatomy of a multi-vector attack: mail bomb + ClickFix

In our Q1 2025 Cyber Threat Report we discussed an attack pattern in which threat actors flooded a target's inbox with junk mail and then followed up via an internal communications platform, posing as the IT help desk and offering to "fix" the problem. That same logic – overwhelm, then rescue – is being adapted with a new payload mechanism. These are the components of the attack:

A mail bomb (or spam bomb) is a saturation attack on an individual mailbox. The threat actor signs the victim's email address up for thousands of newsletters, account confirmations and forum notifications in a short period. That means legitimate alerts, security notifications and routine work emails are quickly buried under hundreds or thousands of messages. The victim is left disoriented, often anxious, and primed to accept help from anyone who appears to offer it.

ClickFix was originally observed as a standalone phishing technique. The threat actor presents the user with what looks like a routine system error or a verification prompt – a CAPTCHA-like page, a reauthentication notice or a "your browser needs to be updated" pop-up. The "fix" instructs the user to copy a string of text and paste it into the Windows Run dialog, into PowerShell or into a terminal. What the user is actually pasting is a malicious command that downloads malware, opens a remote access channel or harvests credentials.

When combined, the sequence reads as a plausible help desk interaction. The mail bomb creates a real, observable problem and a sense of urgency. A vishing call (or seemingly internal message) follows from someone claiming to be IT, offering to walk the user through a fix. The user is directed to a page that completes the ClickFix step. Once the user has pasted the command and pressed Enter, the attacker has the foothold they came for. None of the individual steps look like phishing in the traditional sense: there is no oddly urgent email message, no suspicious link to click, no request to reset a password. The user feels like a participant in their own rescue.

Defending against the mail bomb + ClickFix pattern

The most reliable defense against this pattern is procedural rather than technical. Recommended controls include:

- **Recognize the signature sequence.** An unexpected email flood followed shortly by an inbound "help" contact – whether by phone, internal communication platform or another collaboration tool – should itself be treated as an indicator of attack.
- **Verify inbound help through a known channel.** Employees should never accept the identity of a caller offering IT support at face value. Independently look up the help desk's number or chat handle and reach out from there.
- **Treat "paste this command" as a red flag.** Standard IT support practices do not include asking a user to paste content into the Run dialog, PowerShell or any terminal. This step alone should halt the interaction.
- **Route mail bomb reports to the security operations center (SOC) or IT, not the caller.** When an inbox is suddenly flooded, employees should report the event to the security operations center directly rather than engaging with whoever calls offering to make it stop.
- **Restrict external communications and collaboration messages by default.** As we have noted in prior reports, the default settings on many collaboration platforms allow inbound contact from outside the organization, expanding the attack surface unnecessarily.

Impact Alerts: Q1 2026

Many of the vulnerabilities that were recorded last quarter, including the four highlighted here, share a common issue: authentication and access control mechanism failure. These lead to bypass or circumvention, often without any credentials needed. Some of the most sensitive and widely deployed enterprise tools remain susceptible to pre-authentication exploitation. Travelers issued Threat Alerts to policyholders directly impacted by these vulnerabilities.

A vulnerability in GNU InetUtils telnetd through version 2.7 was published on January 21 and recorded as CVE-2026-24061. It allows a remote attacker to bypass authentication entirely by supplying -f root as the value of the USER environment variable, granting unauthenticated root-level access to any system running the affected telnet daemon. Successful exploitation could enable full system takeover, data exfiltration or use of the host as a launchpad for further attacks into the network. This vulnerability was added to the federal Cybersecurity and Infrastructure Security Agency's (CISA) list of Known Exploited Vulnerabilities (KEV) 5 days after being published. GreyNoise researchers are tracking IP addresses that have been observed attempting to exploit the vulnerability. The vulnerability itself had gone 11 years without being identified.

A vulnerability in multiple Fortinet products, including FortiAnalyzer, FortiManager and FortiOS across several major version branches, was published on January 27 and recorded as CVE-2026-24858. Classified as an authentication bypass via an alternate path or channel (CWE-288), it may allow an attacker with a FortiCloud account and a registered device to log into devices belonging to entirely separate accounts, provided FortiCloud SSO authentication is enabled on those devices. Given how widely Fortinet products are deployed in enterprise and government environments, exploitation of this flaw could allow threat actors to access and reconfigure security infrastructure belonging to other organizations. Fortinet confirmed active exploitation after multiple clients reported that attackers had gained unauthorized access to their FortiGate firewalls and created new local admin accounts. As an immediate mitigation, Fortinet temporarily suspended FortiCloud SSO on January 26 and restored it the following day only on devices that had been patched, with users strongly advised to update firmware to continue using SSO authentication safely.

A critical pre-authentication remote code execution vulnerability in BeyondTrust Remote Support and certain older versions of Privileged Remote Access was published on February 6 and recorded as CVE-2026-1731. By sending specially crafted requests, an unauthenticated remote attacker may be able to execute operating system commands in the context of the site user. BeyondTrust products are frequently deployed to manage privileged remote access sessions across enterprise environments, meaning successful exploitation could hand attackers the keys broader organization access. This vulnerability was added to CISA's KEV a week after publication in February. Researchers at Unit42 have identified campaigns exploiting the vulnerability targeting various sectors including financial services, legal services, healthcare and more in the US, France, Germany, Australia and Canada.

A pre-authentication remote code execution vulnerability in Ivanti Endpoint Manager Mobile (EPMM) was published on January 29 and recorded as CVE-2026-1281. The vulnerability stems from the web server routing specific URL requests to a Bash script that fails to properly sanitize user inputs, allowing attackers to inject malicious payloads via URL parameters and execute arbitrary commands as the web server user. Ivanti confirmed active zero-day exploitation in the wild, and the vulnerability was added to CISA's KEV catalog the same day. Unit 42 observed widespread exploitation including the deployment of dormant backdoors designed to maintain long-term access even after organizations applied patches.

Conclusion

With a record number of ransomware groups active during the quarter, the increasingly fragmented ecosystem will continue to present challenges to organizations seeking to improve their defenses. Meanwhile, the rapid pace of AI adoption inside organizations is itself reshaping the risk picture. Combined with the continued evolution of social engineering into multi-vector campaigns like mail bomb + ClickFix, the trends we observed in the first quarter of 2026 display the need to approach security from multiple angles simultaneously: software governance, training and controls to limit the impact of social engineering and perimeter defense. We hope understanding these trends helps your organization prioritize your security activities in the coming months.

Recommendations from the Travelers Cyber Risk Services Team

To mitigate these risks, organizations should adopt a strong cyber prevention program, including the following recommendations detailing the top security investments with the greatest return on investment.

These recommendations will help increase the bar required for ransomware actors to successfully carry out an attack on an organization.

They include:

- Implement phishing-resistant MFA for all remote access and email.
- Run an effective vulnerability management program to quickly patch critical vulnerabilities in edge devices, such as virtual private networks (VPNs).
- Ensure you have reliable backups and have a resilient disaster recovery and business continuity plan.
- Run endpoint detection and response (EDR) solutions with 24x7 active monitoring.

travelers.com

One Tower Square
Hartford, CT 06183

Travelers analysis was made possible with supporting data from eCrime.ch.