TRAVELERS

Q3 2026

Cyber Threat Report

AI-Aided Exploits Drive New High for Ransomware



Table of Contents

Executive Summary 03

Overview of Global Ransomware Activity in Q3 2026 04

Known Exploited Vulnerabilities Are Climbing, Fueled by Edge Technologies 08

Briefing: What the Luna Moth Leak Confirms About Social Engineering 12

Impact Alerts from Q3 14

Conclusion 15

Published by Travelers with contributions from:

Thomas Nance
Director of Intelligence Services
Darkweb IQ

John Lippe
Director, Cyber Claim Forensics
Travelers

Josh Doguet
Sr. Manager, Incident Response Intelligence
Travelers

Courtney Hassenfeldt
Cybersecurity Technologist
Travelers

Ryan Bell
Director, Threat Intelligence – Cyber Risk Services
Travelers

Alex Pinto
Sr. Director of Product Marketing – Cyber
Travelers

©2026 The Travelers Indemnity Company. All rights reserved.

Executive Summary

It wasn't hard to see coming. Given that the ransomware ecosystem produced record activity in Q4 2025 with minimal abatement in the first two quarters of 2026, and how quickly large language models (LLMs) are being adopted into various cyber-specific use cases, the likelihood of continued increases in ransomware activity seemed almost inevitable.

But "when" and "how" were open questions. We have been watching closely for meaningful impact from AI on cyber risk – something more tangible than hypotheticals or anecdotal reports. Now we're starting to see it. This report finds the most ransomware activity we've ever seen in a quarter, along with credible, widespread usage of AI tooling that makes it easier for threat actors to execute attacks. It appears the new era of risk that many had anticipated has arrived.

But there's plenty of nuance in the details. There are still sources of friction keeping threat actors from being able to leverage AI to drastically increase their activity. Otherwise, we'd be seeing data that looks exponential rather than incremental – and we are not. Probably the single biggest friction comes from the steps taken by organizations to harden their systems. That's the encouraging news: the attacks we're seeing most frequently are still generally preventable with good cyber practices.



A record quarter: Leak sites posted 2,882 victims in Q3, the most we have recorded and 26% more than the previous quarter, with a record 111 distinct groups posting.



Exploited vulnerabilities climb at the edge: The U.S. Cybersecurity and Infrastructure Security Agency (CISA) logged 246 Known Exploited Vulnerabilities (KEVs) to their official list through September 30, already matching the total for all of 2025.

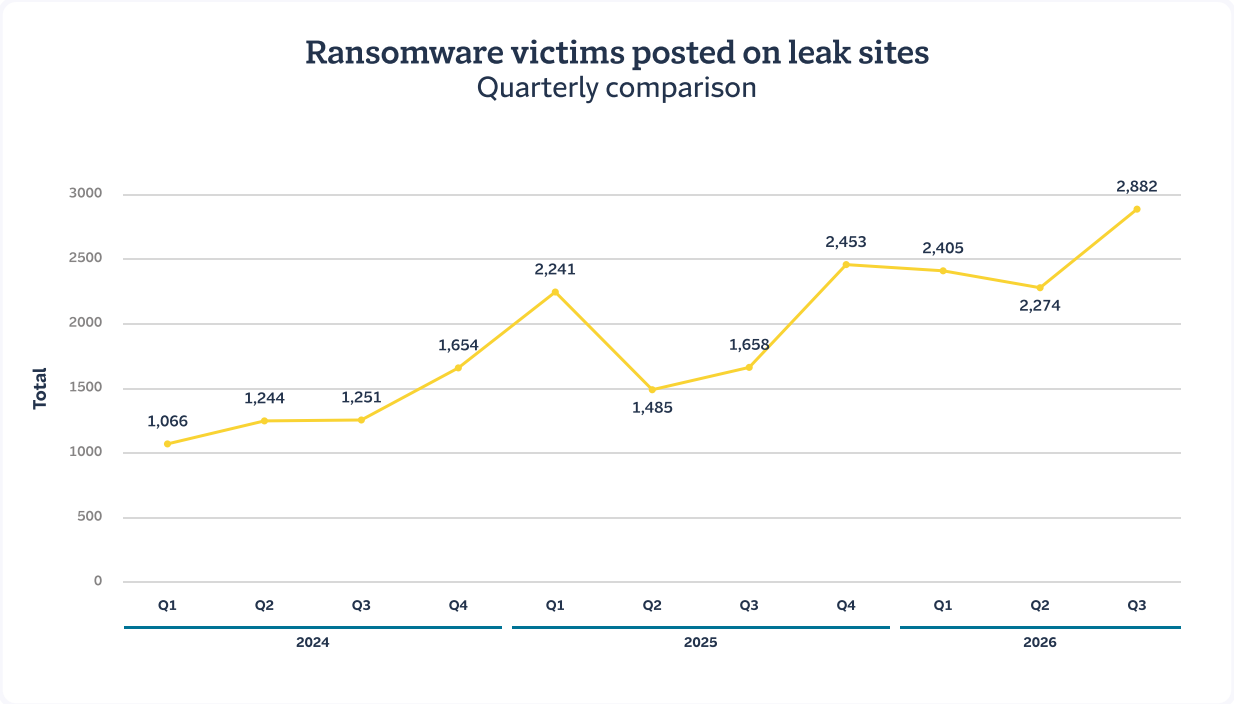


Social engineering drives extortion, too: The Luna Moth leak confirms that phone and email deception, and even in-person agents, are core elements of a strategy to extort businesses.

Overview of Global Ransomware Activity in Q3 2026

Activity reaches a new high

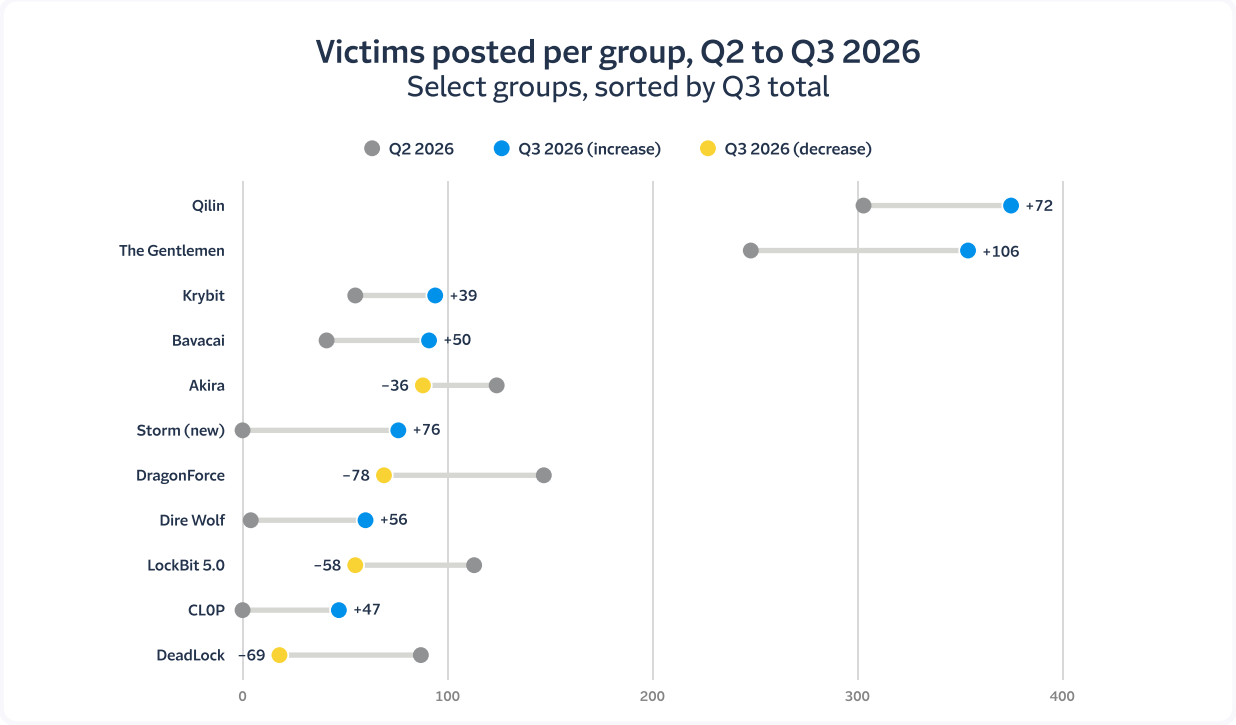
In the third quarter of 2026 we recorded 2,882 victims posted to ransomware leak sites*, the highest quarterly total we have recorded. That is 26% higher than the previous quarter and 429 victims above the previous peak of 2,453 set in Q4 2025. Through the first three quarters of the year, 2026 has produced 7,561 victims, roughly 40% more than the same period in 2025 (5,384).



*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 subset of overall activity, but one that can provide a longitudinal comparison of activity over time.

©2026 The Travelers Indemnity Company. All rights reserved.

August in particular stood out. Our data showed 1,125 victims that month, the busiest of any month we have tracked. Some of the jump in activity can be attributed to new and previously quiet groups like Orova, Storm and Dark Project, none of which had posted before and had 110 victims between them. Those additions came on top of even larger increases in activity from this year's most consistently active threat groups, The Gentlemen and Qilin.



Edge device exploits drive the gain

Some of the spike in August can be traced back to exploit activity that we began to see in our threat intelligence sources weeks earlier.

A prime example was a string of attacks by CL0P, whose all-or-nothing approach we've covered before, against users of a common piece of engineering software in early June. The flaw was patched on June 17 and added to the U.S. government's list of known-exploited vulnerabilities on June 25, yet the victims were not named until August 12, when 42 appeared on the leak site in a single day.

But that vulnerability now looks distinctive for one reason: it wasn't in an edge device (internet-facing equipment like virtual private networks, firewalls and remote-access gateways that sit between internal systems and the open internet). Otherwise, May and June saw a spate of exploits against edge devices.

In early May, an affiliate of Qilin began using a previously unknown flaw in a widely used VPN product to get into victim networks, and roughly a month passed before a fix was published. Qilin was also tied to attacks through a login bypass in a popular remote-access gateway from a different vendor. And in June, a leaked collection of what appeared to be working logins for roughly 74,000 firewalls, built from recycled credentials and weak passwords, arrived along with reports that the data was actively being sold.

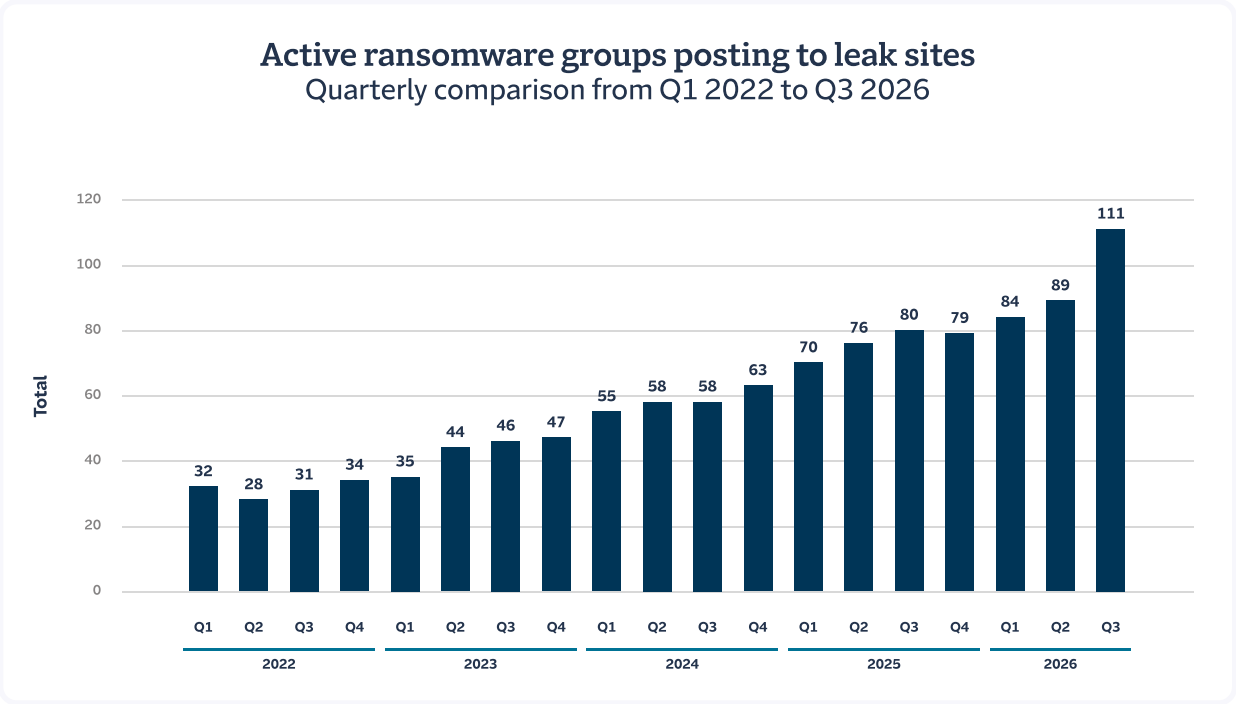
©2026 The Travelers Indemnity Company. All rights reserved.

In retrospect, these edge exploits were the opening salvos in a string of attacks that led to the record-breaking month and quarter a few weeks later. They also represent the continuation of a long-running trend. At the beginning of this year, we identified VPN exploitation as the vector of attack for 44% of the ransomware and system intrusion claims we studied in 2025, up from 36% in 2024. This year's ratio may prove to be even higher.

We'll dig deeper into the topic of edge exploitation in the next section of this report.

Fragmentation, again

Fragmentation and churn remained the defining feature of the landscape for a third straight quarter. A total of 111 distinct groups posted victims, up from a then-high 89 in Q2 and the most in any quarter in our dataset. A record 35 groups posted victims for the first time, accounting for 20% of all victims, while 17 groups that were active in Q2 posted nothing.



It's worth mentioning that "new" means new to our data, not necessarily new to the business. Operators and affiliates often move between brands, seemingly with increasing frequency. Storm, the most active new group, first posted about five weeks after a similarly named group, Stormous, went quiet. In cases like that it's easy to make some assumptions about what the new group's characteristics might be. But other times the personnel movements are more opaque. The challenge for defenders is that more groups mean more different combinations of tactics and skill sets, raising unpredictability.

That matters for law enforcement, too. An international coalition of authorities seized the KillSec group's leak site on September 30. But with ransomware spread across this many groups, a takedown has less impact than it did in the past when one enforcement operation could dampen the entire ransomware ecosystem for weeks or months.

©2026 The Travelers Indemnity Company. All rights reserved.

Industry spotlight: law firms

Law firms are attractive targets for a simple reason: they hold other people's secrets. Client strategy, deal terms and litigation files are the kind of material a firm would pay a great deal to keep private, which gives an extortionist leverage without locking a single computer.

The FBI warned in May that Luna Moth, also known as Silent Ransom Group, was stealing data from law firms through social engineering and impersonation, then demanding millions of dollars to keep it from being published. A recent leak of the group's files, which we cover in our Briefing later in this report, corroborated that warning. There are repeated references to targeting the world's top 100 law firms, and researchers have counted 42 law firms among the group's 50 most recently observed victims.

The leak also helps to confirm observations made by our Claim professionals in Q3, who noted rising instances of law firms being targeted as well as a sharp rise in the proportion choosing to pay a ransom. For firms, the most useful defense is procedural: staff should always verify any unexpected IT or support request and do so using a form of contact they already know, not one supplied by the requester.



There are repeated references to targeting the world's top 100 law firms, and researchers have counted 42 law firms among the group's 50 most recently observed victims.

©2026 The Travelers Indemnity Company. All rights reserved.

Featured Analysis: Known Exploited Vulnerabilities Climb with AI Tooling

Let's get some acronyms out of the way: the Cybersecurity and Infrastructure Security Agency (CISA) maintains a public list of Known Exploited Vulnerabilities (KEVs). This is a complement to the Common Vulnerabilities and Exposures (CVE) catalog that we reference frequently. Our team uses both lists in our research on threats that may affect policyholders.

The difference is that the KEV list only includes software flaws that are confirmed to have been used in real-world attacks, whereas the CVE list includes any that has been discovered and reported. The KEV list is shorter by orders of magnitude: less than two thousand entries against more than 250,000 CVEs.

We were curious about the rate of change in KEVs versus CVEs because of the impact we anticipate AI tools are having on vulnerability discovery. We have already seen the pace of CVE additions start to grow, but would KEVs increase in kind? In other words, will attack activity correspond to the increase in vulnerability discovery, or are the defenders maintaining an upper hand?

KEVs: not exploding, but increasing

In 2026, the catalog has grown faster than in any prior year, and the gap between a flaw becoming public and its appearing on the list has narrowed to a matter of days.

By the numbers — CISA KEV catalog, 2026 year-to-date:

246
vulnerabilties
added through September 30, already matching all of 2025 (245).

50
edge device CVEs
published in 2026 have reached the catalog, up 35% from 37 in all of 2025 — the highest total of any publication year.

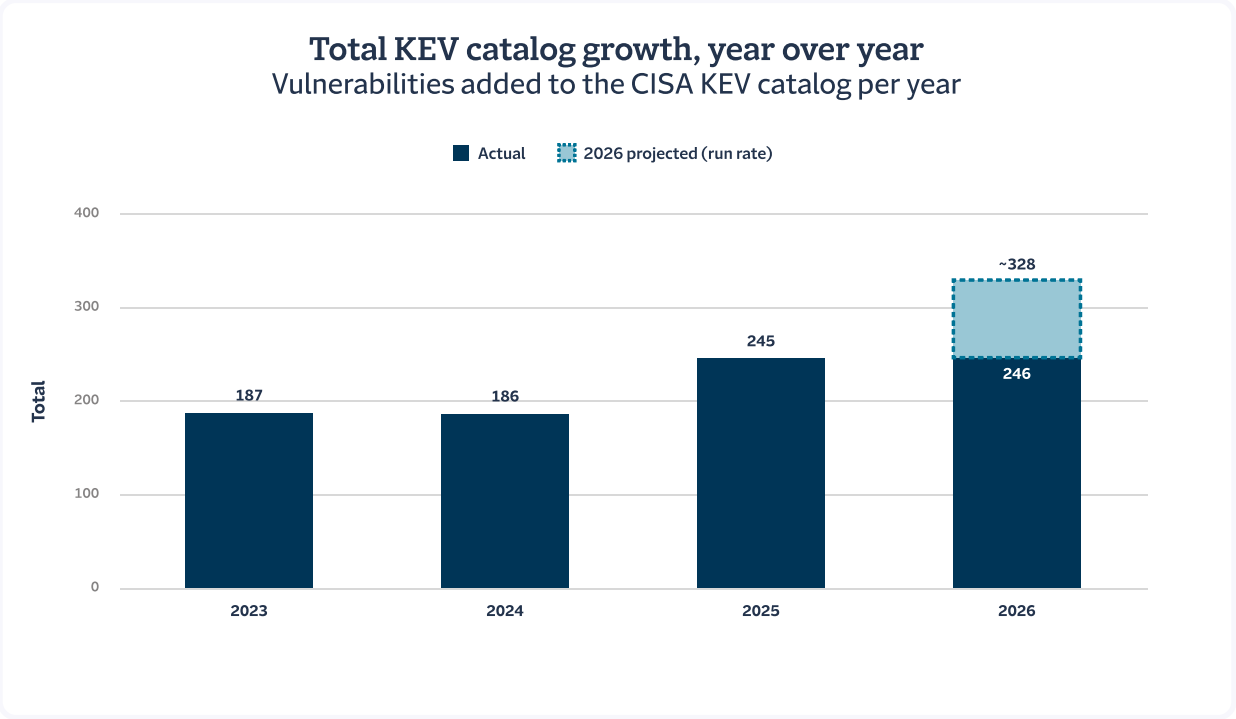
About
28%
of 2026's newly listed CVEs are edge devices, up from 19–23% across 2022–2025.

14
days:
the median gap from a 2026 CVE's publication to its listing.

By the end of September, CISA had added 246 vulnerabilities to the catalog in 2026, already exceeding the 245 added in 2025 with a quarter still to go. A simple run-rate projection would put the 2026 total near 329 additions, roughly a third more than 2025. This aligns with growth of roughly 30% in the previous year.

©2026 The Travelers Indemnity Company. All rights reserved.

By the end of September, CISA had added 246 vulnerabilities to the catalog in 2026, already exceeding the 245 added in 2025 with a quarter still to go. A simple run-rate projection would put the 2026 total near 329 additions, roughly a third more than 2025. This aligns with growth of roughly 30% in the previous year.



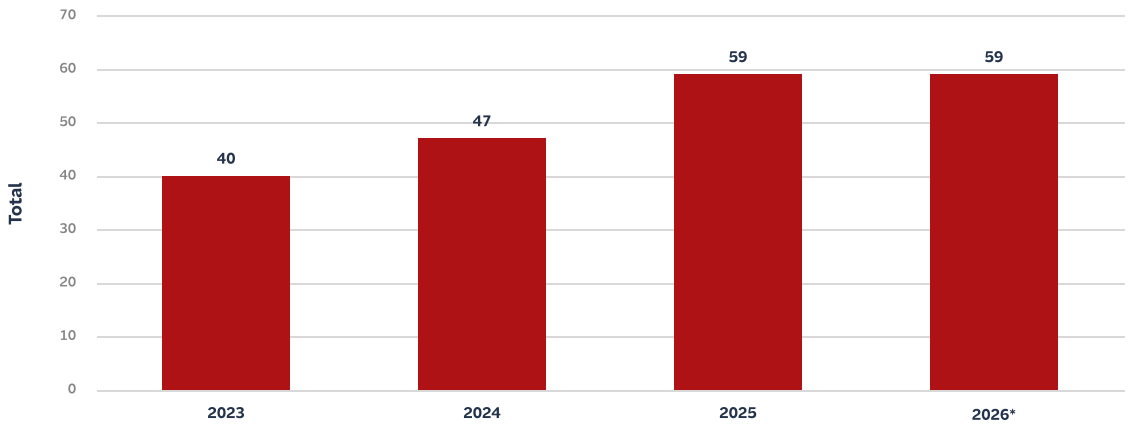
These figures certainly don't point to any kind of drastic, overnight change in the number of vulnerabilities being exploited. But consistent growth and the compounding effect of many simultaneous active KEVs certainly impact how organizations approach defense.

©2026 The Travelers Indemnity Company. All rights reserved.

Risk increasingly starts at the edge

The clearest grouping we see within the list of KEVs is of edge and perimeter devices — the VPN gateways, firewalls and remote-access appliances that sit between the public internet and the internal network. Fifty-nine edge device vulnerabilities had reached the KEV catalog through the end of September, matching the number for all of 2025.

Edge device vulnerabilities added to the KEV catalog
By year CISA cataloged them, 2023 to Q3 2026



*2026 includes data from Q1 to Q3 2026.

This higher volume of edge KEVs also represents a larger share of the overall entries, at roughly 28% of 2026-published entries, up from the 19–23% range that held from 2022 through 2025. The third quarter was especially heavy: September alone brought 16 edge KEV additions, more than a quarter of the 59 recorded so far this year.

In our reports covering Q3 and Q4 of last year, the main topic was a coordinated campaign by several threat actor groups to target a critical VPN vulnerability. The KEV data from this year contextualizes that campaign as a key event within a steadily expanding strategy, rather than an opportunistic aberration.

AI is widening the field

One takeaway from this data is that we're not yet seeing the nightmare scenario of fully automated attack chains, which would likely yield an explosion in attacks. Nonetheless, the increase in KEVs is pronounced and consistent. It's worth exploring how AI is potentially impacting the threat ecosystem to yield these outcomes.

We believe a major factor is the broadening of the pool of individuals who can turn a known vulnerability into an exploit.

©2026 The Travelers Indemnity Company. All rights reserved.

The fact that there are a quarter-million CVEs but fewer than 2,000 KEVs shows that exploits aren't a simple matter. Specialized skill and effort are required to develop an attack, and the payoff for investing in something new is far from certain. (And it's not just a matter of time, either: many CVEs have been published for years — decades, even — but haven't been exploited.) That's why we see threat actors gravitate to playbooks that are proven to work and easy to copy. It's also where AI's rapidly improving ability to code and test software can make an impact.

Threat intelligence from Darkweb IQ, drawing from undercover engagements with criminal sellers, describes LLM assisted exploit development spreading across the initial access broker (IAB) market. Traditionally, IABs are individuals or groups who specialize in breaking into networks and reselling access to others who will exploit it through ransomware, data theft or fraud. But recently, Darkweb IQ found IABs using LLMs to help build exploit code to offer along with the access they're selling.

One IAB tracked by Darkweb IQ since late 2024 began as a credential and log checker, selling access harvested from a type of malware called an infostealer. By 2026, though, the IAB was supplying functional exploits for edge device flaws that were assembled with the help of a coding agent following public disclosures of the vulnerabilities. Another seller with a similar profile started offering a bundle of exploits for six vulnerabilities. That actor, in Darkweb IQ's assessment, was not offering any new capabilities — just a more refined path to carry out attacks.

"We are seeing more actors use LLMs to develop working exploits, including actors who started out brute forcing VPNs and selling infostealer credentials" said Thomas Nance, Director of Intelligence Services at Darkweb IQ, "and we are observing actors moving into more complex initial access pathways than they had demonstrated before."

> "What they're exploiting isn't new. These are the same known, patchable weaknesses actors have exploited for years. But the LLM took them further and faster than they would have gotten on their own."

This assessment was echoed in Anthropic's September 2026 account of how financially motivated actors used its models. None of the documented operations depended on a technique genuinely novel to defenders — initial access still ran through unpatched edge devices, stolen credentials and public vulnerability research. As Nance puts it, "What they're exploiting isn't new. These are the same known, patchable weaknesses actors have exploited for years. But the LLM took them further and faster than they would have gotten on their own."

©2026 The Travelers Indemnity Company. All rights reserved.

What it means for defenders

When we see, like we did this quarter, a record number of distinct groups operating and increasing churn in the ecosystem, there's an easy connection to the kind of activity observed by Darkweb IQ and Anthropic. More actors are capable of developing workable exploits. That leads to more and better options available "off the shelf" which, in turn, leads to more actors striking out on their own.

That said, AI isn't a magic wand. Even with LLMs, "Not everyone can develop exploits — at least not yet," said Nance. "Getting working exploit code out of a model takes skill, and the actors that are succeeding have that skill." Thus the steady, but not explosive, increase in KEVs.

Considering all of this, our guidance for a defensive response is both unchanged in substance and more urgent. The best-positioned organizations are those that can find, prioritize and patch known-exploited vulnerabilities on their perimeter faster than the shrinking window allows.

For Travelers policyholders, we strongly recommend acting quickly upon alerts you might receive from us or from your cybersecurity vendors. When a new vulnerability is published (particularly if it is added to the KEV list), we use our scan technology to identify policyholders using the affected device or software. That means a policyholder will only be alerted when a vulnerability has direct relevance to their organization, with only occasional exceptions for very widespread web technologies. Each alert includes remediation steps to help the policyholder quickly address the vulnerability.

Briefing: What the Luna Moth leak confirms about current attack strategies

Most of what we know about how a ransomware group operates is inferred from the outside — from leak-site postings, malware samples, victim telemetry and cryptocurrency flows. From these we can make assumptions about operational models and tactics that are solid enough to help build defensive strategies and help our policyholders. But occasionally, we get the benefit of more than just a peek.

The leak of key files, chat exchanges or financial data can reveal how a group actually operates, as with the leaked IAB training manual we covered back in 2024, or the internal chats leaked from Conti in 2022 and Black Basta in early 2025. One of the most detailed leaks to date came to light recently: that of a group known variously as Luna Moth and Silent Ransom Group.

Inside the Luna Moth leak

One assumption that has been hard to pin down from the outside is how deliberately certain groups rely on social engineering — tricking an employee into granting access — rather than on purely technical intrusion. The Travelers Claim team has recently observed more cases that appeared to include the use of social engineering with the specific aim of enabling a ransomware attack. But were those cases based on pure opportunism?

©2026 The Travelers Indemnity Company. All rights reserved.

Briefing: What the Luna Moth leak confirms about current attack strategies (continued)

The recent leak, which included internal chat logs, ransom records and cryptocurrency wallet addresses, confirms that social engineering is not merely one of the group's methods but in fact represents the core of its operation. As documented in the internal brief, Luna Moth operators deceive staff into granting access through fake-invoice emails and "IT support" phone calls, and since 2025 have escalated to recruiting local agents in target locations and sending them into offices in person using fake identification.

The group formed in early 2022 as the first faction to break away from the Conti ecosystem, and it built its operation around "callback" phishing — a technique in which a lure email prompts the victim to call a phone number, where an operator talks them into granting remote access. The recent leaks show that this genre of email- and phone-based social engineering has not diminished.

Following the money

Our team has also observed pieces of how these groups move and launder their proceeds, but rarely with the whole picture in view. The leaked chats fill in the gaps, at least for this particular web of threat actors.

According to Crystal Intelligence's analysis, the group's operations lead set strict rules for receiving payouts: use a fresh wallet for every payment, never reuse one flagged as "unclean" and keep each victim's money separate. The operators even screened their own coins before spending them, then converted them to cash through instant exchangers paying out to foreign bank cards, a cash courier in Eastern Europe and a service that swapped Bitcoin for payments via digital wallets.

The discipline did not always hold. In one case, a recruiter spent two payouts together, linking wallets the rules were meant to keep apart, and smaller payments went to regulated exchanges, where accounts are tied to verified identities. Tracing often succeeds for just this reason: human error rather than technical weakness. The leak has helped confirm the rest of the pattern, which was inscrutable by design.

Law firms in the crosshairs

The leak also shows who the group was after, and how far it would go to reach them. About 50 organizations appear in its records, most of them law firms, which fits the pattern we describe in our spotlight on law firms earlier in this report.

To get past the front door, the group did not rely on phone calls alone. According to the chat logs, actors directed U.S.-based "agents," recruited through paid social media advertisements disguised as ordinary job listings for courier services, nightclub promotion and security.

In one negotiation, a law firm's representative told the group the firm knew that an individual had entered its New York office and copied files onto a flash drive. The FBI had separately warned earlier this year that members were posing as IT personnel to gain physical access to computers.

The group's internal records claimed roughly $207 million paid by 27 firms over six months. Independent analysis by Chainalysis did not confirm those exact figures but did validate that the leaked wallet addresses were authentic and that large sums moved roughly as the chats described, including one wallet funded entirely by a confirmed victim payment exceeding $10 million.

©2026 The Travelers Indemnity Company. All rights reserved.

Briefing: What the Luna Moth leak confirms about current attack strategies (continued)

Why this matters

Our claim trends show a similar pattern of law firm victims, with very sensitive data exfiltrated during the attacks and a preference to pay in exchange for the deletion of the client data. Luna Moth has continued using this successful playbook in October with several law firms posted on their data leak site. Confirming the details of the group's strategies helps our team to sharpen our incident response tactics and the advice we give to policyholders that are affected. It also helps inform the work of our Cyber Risk Services team, who work with policyholders every day through consultative sessions including the Social Engineering Fraud risk assessment, which was introduced earlier this year. (Travelers Cyber policyholders can contact the team any time at cyberservices@travelers.com to discuss availability of services.)

Impact Alerts from Q3

A selection of vulnerabilities gathered by the Travelers Threat Intelligence team. Travelers issued Threat Alerts to policyholders directly impacted by these vulnerabilities.

CVE	Technology Affected	Type of Vulnerability	Impact
CVE-2026-18577	N-able N-central (remote monitoring and management platform); builds prior to the 2026.3	Authentication bypass and account takeover	This was the result of an incomplete patch for CVE-2026-18556, which was exploited in the wild as a zero-day and added to the CISA KEV catalog on August 3, 2026. Microsoft Threat Intelligence reported that the group Storm-1175 is suspected of exploiting it to deploy a new ransomware family, StormEncryptor.
CVE-2026-88771 **CVE-2026-88772**	Citrix NetScaler ADC and NetScaler Gateway (Citrix bulletin CTX697096)	**CVE-2026-88771:** unauthenticated command injection leading to remote code execution **CVE-2026-88772:** memory overflow leading to remote code execution or denial of service	Both were exploited as zero-days to plant webshells and added to the CISA KEV catalog on September 27, 2026, with a three-day remediation deadline for federal agencies (September 30). Initial CVE-2026-88772 intrusions were attributed to suspected state-sponsored actors; no specific ransomware group was named in the sources we reviewed.
CVE-2026-94127	F5 BIG-IP Access Policy Manager (APM) configured with an OAuth profile on a virtual server (OAuth authorization server)	Heap-based buffer overflow enabling unauthenticated remote code execution	F5 confirmed exploitation when it published its advisory on September 22, 2026, and CISA added it to the KEV catalog the same day, with a three-day remediation deadline for federal agencies.

©2026 The Travelers Indemnity Company. All rights reserved.

Conclusion

This quarter's data shows a clear story of steady growth in attack activity. AI is part of that story, because it lowers the skill needed to turn a known flaw into a working exploit and widens the pool of actors who can run one. Yet the solutions are not so cutting edge. The Luna Moth leak shows that some of the most effective attacks rely on little more than a convincing phone call or email, while the exploits behind the edge device surge target flaws that have had patches available for weeks or months. The organizations best positioned for what comes next will close the window quickly on known-exploited vulnerabilities at the perimeter, protect remote access with phishing-resistant MFA and train their teams to verify unexpected requests before acting on them.

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.

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 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.

Built for cyber.

With always-on threat intelligence, we're able to help brokers and policyholders outpace cyberattacks.

Learn More

©2026 The Travelers Indemnity Company. All rights reserved.



travelers.com

One Tower Square
Hartford, CT 06183

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

Travelers Casualty and Surety Company of America and its property casualty affiliates. One Tower Square, Hartford, CT 06183.

This material is for general informational purposes only and is not legal advice. It is not designed to be comprehensive and it may not apply to your particular facts and circumstances. Consult as needed with your own attorney or other professional advisor. This material does not amend, or otherwise affect, the provisions or coverages of any insurance policy or bond issued by Travelers. It is not a representation that coverage does or does not exist for any particular claim or loss under any such policy or bond. Coverage depends on the facts and circumstances involved in the claim or loss, all applicable policy or bond provisions, and any applicable law. Availability of coverage referenced in this document can depend on underwriting qualifications and state regulations.

Cyber customers may receive certain services through external vendors and, if using these services, must agree to the vendors' terms of use and privacy policies. Travelers makes no warranty, guarantee or representation as to the accuracy or sufficiency of any such services. The use of such services and the implementation of any product or practices suggested by such vendors is at the customer's sole discretion. Travelers disclaims all warranties, express or implied. In no event will Travelers be liable in contract or in tort for any loss arising out of the use of such services or any vendor products. Claims scenarios are based on actual claims, composites of actual claims, or hypothetical situations. Resolution amounts are approximations of both actual and anticipated losses and defense costs. Facts may have been changed to protect confidentiality.

© 2026 The Travelers Indemnity Company. All rights reserved. Travelers and the Travelers Umbrella logo are registered trademarks of The Travelers Indemnity Company in the U.S. and other countries. BSI-CY-0143. New 11/26.