TRAVELERS

Q2 2026

Cyber Threat Report

Account Takeovers Rise

With an assist from AI

Table of Contents

Published by Travelers with contributions from:

Carolyn Purwin Ryan
Partner
Mullen Coughlin

Josh Doguet
Sr. Manager, Incident Response Intelligence
Travelers

Brad Roughan
AVP, Threat Intelligence
Travelers

John Lippe
Director, Cyber Claim Forensics
Travelers

Courtney Bergeron
Cybersecurity Technologist
Travelers

Alex Pinto
Sr. Director of Product Marketing – Cyber
Travelers

Executive Summary

Ransomware activity eased slightly in the second quarter of 2026 but stayed well above historical norms — 2,274 victims were posted to leak sites, down about 5% from Q1 but up 53% year-over-year, and the third-highest quarterly total we have tracked. The field also kept fragmenting, with 89 distinct threat actor groups active during the quarter, a new high in our dataset. The more consequential movement this quarter, though, appears elsewhere in our claims data.

Account takeovers surged. Business email compromise (BEC) claims ran 57% higher in the first half of 2026 than in the same period a year earlier. Much of that growth coincides with a shift in tactics: attackers are increasingly stealing authentication tokens rather than passwords, using phishing kits that are augmented by AI tools to take over a victim's full enterprise workspace.

Another area of rising exposure has little to do with threat actors and breaches. Litigation over website tracking technologies — pixels, cookies and session-replay tools — continued to accelerate. Even when setting aside filings from a single serial plaintiff whose activity reached a peak in Q2, pixel-related claims at Travelers across the first half of 2026 ran nearly 60% higher than the 2025 quarterly average. Once concentrated in healthcare, this exposure is now spread broadly across the economy, driven by plaintiffs applying decades-old privacy statutes to routine web tracking. We examine that litigation with the help of a leading expert later in the report.

Activity remains elevated: 2,274 victims were posted to leak sites in Q2 2026 — down about 5% from Q1 but up 53% year-over-year, and the third-highest quarter we have tracked.

Account takeovers rise: BEC claims were 57% higher in the first half of 2026 than the same period in 2025, with a notable trend toward the use of session tokens to bypass authentication controls.

And cyber litigation is rising too: Claims continue to arise from plaintiffs alleging violations of the California Invasion of Privacy Act (CIPA)

Ransomware declines (modestly)

The second quarter of 2026 recorded 2,274 victims posted to ransomware leak sites*, a modest step down from the near-record levels of the previous two quarters. While this represents a decline of about 5% from Q1 2026's 2,405 victims, the quarter tells a familiar story when you consider that this level of activity would have looked extraordinary a year ago. It's now the norm. A direct comparison reveals that Q2 2026 activity was up 53% against the same quarter last year.

The quarter-over-quarter decline deserves examination, though. In 2025, the second quarter saw a steep seasonal decline, falling nearly 34% from Q1. This year, the comparable drop was less than 5%. Whatever seasonal forces have historically pulled second-quarter activity down appear to have been largely overwhelmed by the sustained elevation in overall activity. As we noted last quarter, each new high watermark increasingly tends to set the floor for what follows rather than a temporary peak.

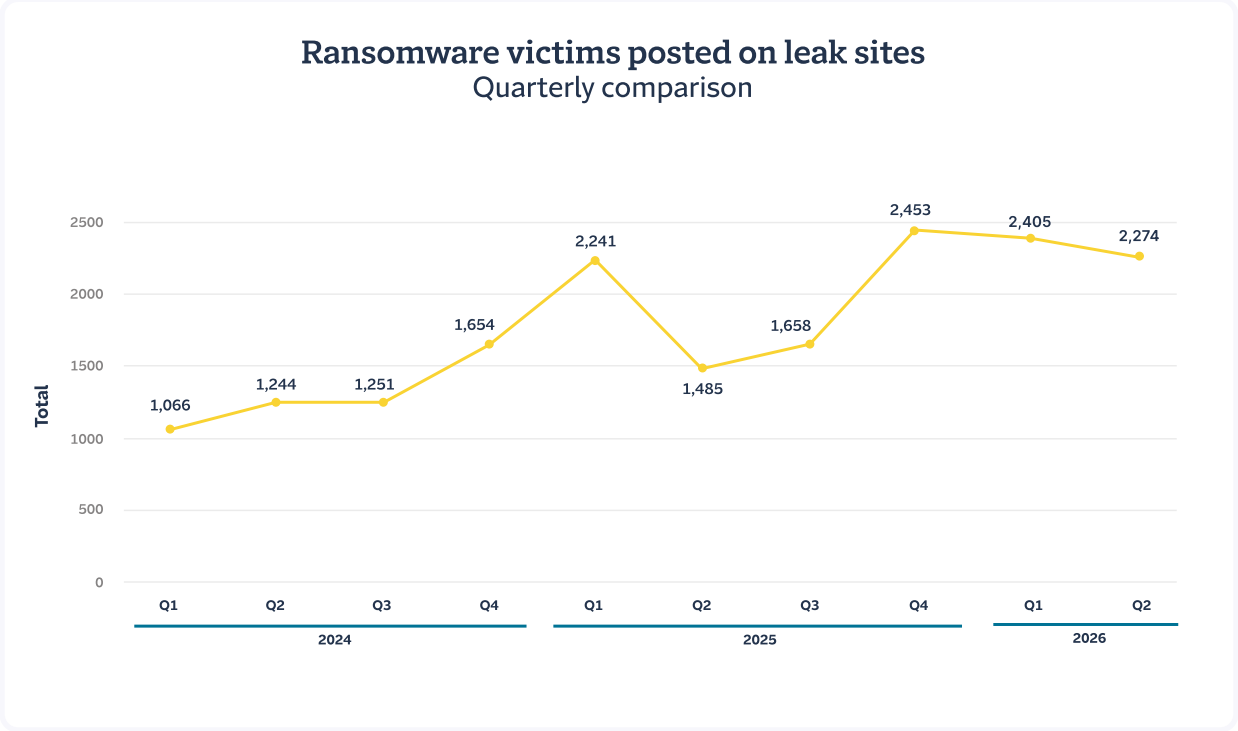

*Data shared on leak sites provides a proxy for overall ransomware activity. A victim's information will typically be posted if the victim has refused to pay a ransom. This means that the data should be viewed as a subset of overall activity, but one that can provide a longitudinal comparison of activity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

For a second quarter, fragmentation defines the landscape

Last quarter we described fragmentation as "the word of the quarter." In Q2 2026, that trend not only persisted but intensified. Eighty-nine distinct threat actor groups were active during the quarter — a new high in our dataset, up from 84 in Q1 2026 and 76 in Q2 2025. Twenty of those groups appeared in leak site data for the first time.

The churn that we highlighted last quarter also shows no sign of slowing. In Q2 2026, 21 groups that had been active in Q1 went quiet, while 20 new ones debuted. This constant turnover — with roughly as many groups exiting as entering each quarter — has become a structural feature of the ecosystem rather than a temporary phase. The result is a threat environment that is not only larger, but harder to map, as the roster of active operators is meaningfully different from one quarter to the next.

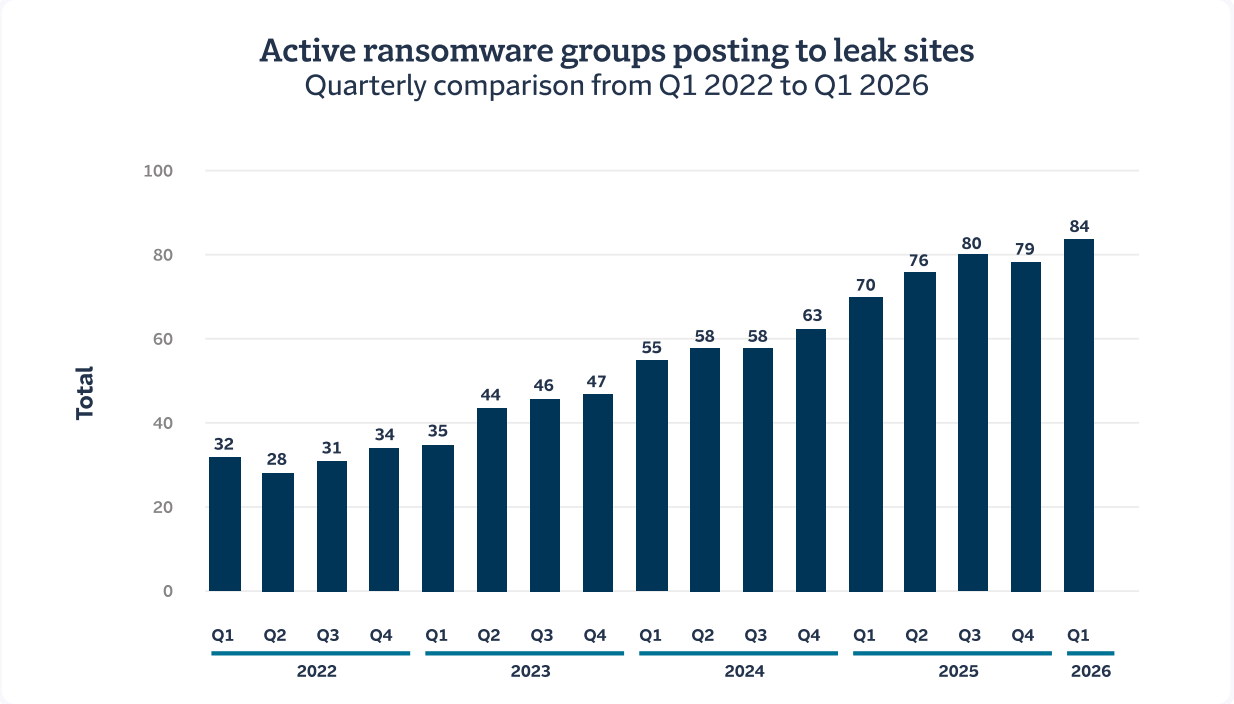

Consistent with a more crowded field, the share of activity captured by the largest groups continued to erode. The three most active groups accounted for about 31% of all leak site postings in Q2 2026, down from 34% in Q1 2026 and 37% a year earlier. The gradual decline in concentration reflects that attacks are being distributed across an ever-larger population of operators, which makes the possibility of disrupting of any single group through law enforcement activity less consequential to the overall volume of activity.

> Eighty-nine distinct threat actor groups were active in Q2 2026 — a new high in our dataset — continuing the fragmentation trend we identified last quarter.

Top groups remain active...

Qilin remained the single most active group for the third consecutive quarter, posting 303 victims. But that figure represents a notable pullback from the 414 victims it posted in Q1 2026 and its Q4 2025 peak of 551.

> A group that did not exist in our data a year ago is now within striking distance of the most active ransomware operation in the world by this measure.

The reason is the continued surge of the group operating as "The Gentlemen." After appearing in September 2025 and posting 207 victims in the first quarter this year, the group was still more active in Q2 with 248 victims. A group that did not exist in our data a year ago is now within striking distance of the most active ransomware operation in the world by this measure. (For background on The Gentlemen's unusually rapid rise, see our Q1 2026 report.)

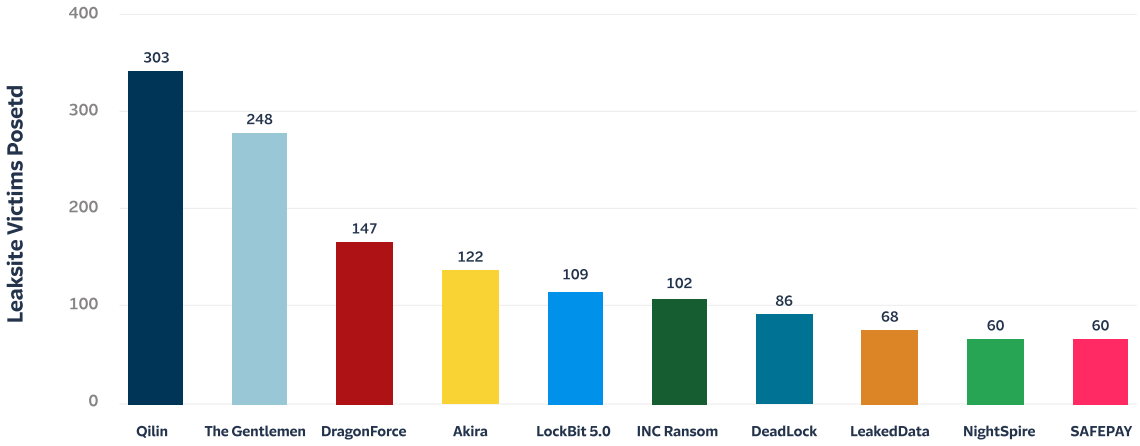

...And new entrants keep arriving

Among the 20 groups making their first appearance in Q2 2026, the most active was a group with the calling card “Bavacai” (among other similar name variants), which posted 51 victims — enough to rank among the 20 most active groups overall. Their sudden rise in activity echoes the pattern we saw with The Gentlemen in recent quarters. New operators are increasingly capable of reaching significant scale immediately, rather than building up gradually over several quarters.

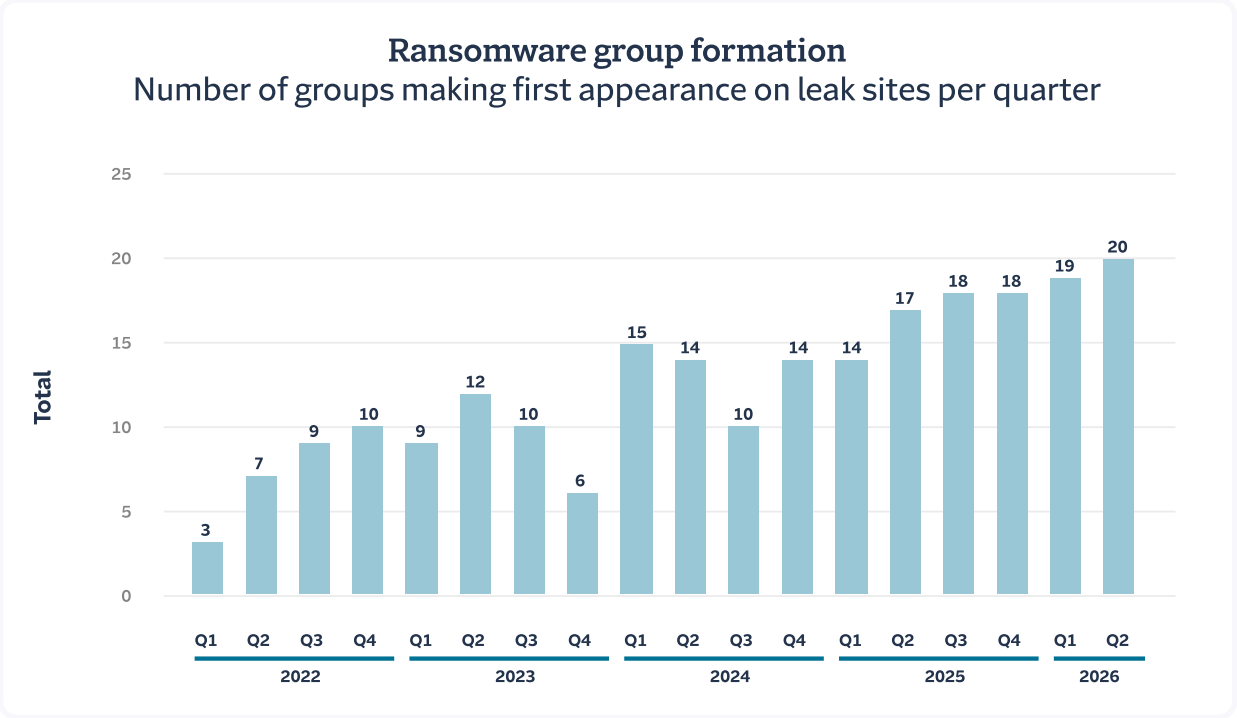

Or perhaps it’s more accurate to say “new” operators. With the maturation of the ransomware-as-a-service (RaaS) model, a new group name does not necessarily indicate a new or inexperienced operator. Affiliates and operators move between brands, tools and infrastructure available for acquisition, and a group can stand up a high-volume operation in a matter of weeks. (The Bavacai group, for instance, has been linked by researchers to the MedusaLocker RaaS operation, seeming to be an offshoot with distinct tactics and technology.) For defenders, this means that tracking a fixed list of “groups to watch” is of limited value when a previously unknown name can become a top operator in a single quarter.

Revisiting Mythos: What Project Glasswing Has Revealed

In Cyber Defense in the World of Mythos, an article published in June, we described a potential step change in cybersecurity. For years, we wrote, the advantage of scale has belonged to attackers – and now frontier AI models have the potential to supercharge that dynamic.

The model at the center of the concern is Claude Mythos, which its developer, Anthropic, judged to be so capable at finding and exploiting software flaws that it withheld its general release. Instead, the company granted a vetted group of organizations access to the model in what it called Project Glasswing, allowing participants to discover their own software's vulnerabilities (and fix them) before the model could be used against them. In the article, we flagged three ways the status quo could shift if Mythos-like capability became available to attackers:

1. Vulnerabilities could be found by threat actors that were previously unknown to researchers or security teams, leading to an increase in "zero-day" exploits (those in which exploits occur before a patch has been issued to fix the vulnerability).

2. Threat actors could use automation to bring down the cost, in time and resources, of exploiting certain types of vulnerabilities, leading to an increase in attacks on known vulnerabilities that didn't make economic sense before.

3. When vulnerabilities are discovered by researchers and published on the Common Vulnerabilities and Exposures (CVE) database, the window of time between publication and active exploitation could further shrink, giving organizations less time to react when information is published.

Glasswing's proof in the pudding

In the weeks since then, Glasswing participants have begun publishing the results of their experimentation with the model. We're beginning to see how likely any of those three outcomes may be.

Mozilla, the maker of the Firefox browser, fixed 271 vulnerabilities in a single release – roughly ten times what the same team surfaced a cycle earlier with a general-purpose model. Cloudflare ran the model against dozens of its own repositories and found it could chain separate flaws together into working proof-of-concept exploits. Palo Alto Networks reported accomplishing a year's worth of penetration testing in under three weeks.

One thing that's clear from the reports: predictions about a coming collapse in the cost of vulnerability discovery are bearing out. And if uncovering vulnerabilities is no longer the primary challenge, security will increasingly depend on how quickly an organization can verify, prioritize and remediate them. Anthropic said in its own summary of the project results that "Currently, there's often a long lag between the discovery of a vulnerability...and the time when the patch is widely deployed by end users...Mythos-class models significantly shrink the time and cost required to find and exploit vulnerabilities, magnifying the risk associated with these time lags."

Anthropic's conclusion supports the idea that the third possibility above is the most likely to manifest soon. Already, one estimate finds that the median time to exploitation of new vulnerabilities has fallen from about a year in 2021 to just a single day this year. Mythos will likely compress that further.

That's particularly concerning given that the average time to patch vulnerabilities has been trending up, not down. A recent publication from J.P. Morgan reported that the average time to patch vulnerabilities climbed from around 60 days in 2023 to nearly 100 days this year. And the issue of speed could be compounded by sheer volume, too. As Mythos and models like it become more widely used, the discoveries published in the CVE database could rise in number – imagine mapping Mozilla's purported 10x increase in patching across thousands of pieces of software. Keeping up is going to require much more effort on the part of software users, many of whom do not have established patching programs as it is.

Don't despair – but don't delay

One thing hasn't changed since our original article was published in June: Travelers' own claims and incident data still show no sign that attackers have gained significant new capability in vulnerability discovery. That is why our recommendation remains the same, but has grown more urgent: treat patch timelines as hard operational deadlines, not best-effort targets. The window between a flaw becoming known and becoming exploited is the one thing this technology is certain to compress.

An in-depth set of recommendations for increasing patching speed and defending against breaches is available in our article: Cyber Defense in the World of Mythos.

Account Takeovers Grow Amid Shift to Token Theft Tactics

Cat-and-mouse game: a (very) brief history of BEC

The earliest approach to compromising accounts, dating back decades, was to steal or brute-force a password. We all experienced the escalating degrees of password complexity we were asked to adopt as a result – but ultimately no amount of characters can defend against your password being harvested and sold on the dark web.

What finally worked to undercut password compromises was the wide adoption of multifactor authentication (MFA). MFA takes away the path of least resistance for attackers and has become standard for almost any online account where sensitive data might be stored.

Attackers then adapted to perform adversary-in-the-middle (AiTM) attacks, in which the victim is tricked into using a fake login page controlled by the attacker, or where the attacker intercepts the secret MFA code by some other means. We have covered AiTM in prior reports, and it remains active. But more organizations are now using phishing-resistant forms of MFA, which involve requiring the use of a passkey or device certificate instead of a secret code, reducing the risk of the most common forms of AiTM attacks.

Now, the latest evolution of attacks sidesteps the login page entirely: token theft. Organizations should adopt defenses against this kind of attack, just as they have in the past, and retain the upper hand in this game of cat-and-mouse.

As the previous section makes clear, there is a widespread assumption that threat actors will – at some point – begin to use AI to exploit software vulnerabilities at a significantly larger scale. While we wait for that eventuality, we've observed AI cropping up in other ways. We have seen it in the use of deepfaked videos and phone calls in social engineering attacks (covered in our Q1 2025 report). We also noted an interesting use of AI in the process of ransomware negotiations (Q1 2026 – see sidebar). While those situations are novel, they have not significantly driven claim trends. That may now be changing with a surge in account takeover activity in 2026.

Business email compromise surges

For years the insurance industry has used "business email compromise" (BEC) as shorthand for when an attacker gains access to – or credibly impersonates – a corporate email account. But that term undersells what we're now seeing in many of these claims. Today, a BEC incident typically encompasses an entire online work identity: for example, a Microsoft 365 or Google Workspace account that includes email, file storage, collaboration tools and cloud applications. These full account takeovers offer more avenues for threat actors to gain information and extract money than email access alone.

The ability to get more out of a single breach is perhaps one factor explaining the significant rise in the frequency of these kinds of attacks. Through the first half of 2026, the average number of BEC claims per month was 50% higher than the monthly average across 2025. Isolating the same period a year earlier makes the shift more pronounced: the total number of BEC claims in the first half of 2026 was 57% higher than the total during the same period in 2025. By contrast, ransomware claims have seen no corresponding rise this year.

But the trend toward universal workspace accounts has been happening for several years. The target has been there. So why the sudden spike now?

Token theft and device code phishing

A new form of account takeover has been one of the primary drivers of the spike in frequency.

When a user signs into an online account, the application is typically configured to hand their device a digital token that spares them from having to enter their password again for some period of time. Attackers have found a new way to acquire that token, meaning they don't need to contend with any authentication controls, whether passwords or other factors in a multifactor authentication (MFA) setup. The token is proof that authentication has happened, so the attacker, with token in hand, bypasses it entirely.

The Travelers team has been tracking the use of an AI-enabled phishing kit known to researchers as Kali365, which has been observed since April 2026. With this kit the attacker has access to an online panel that appears to have been "vibe-coded" with AI tools, from which they send the victim a phishing email designed to look like it's from a Microsoft app. That email directs them to Microsoft's real sign-in page and includes a short code for the victim to type into the page. (This works because the attacker has exploited an authorization flow – OAuth – built for devices like TVs or printers – think of how you can easily log into a streaming service on a new TV via a short code linked to your smartphone's app.)

The victim signs in on the genuine Microsoft site, and everything looks normal — but the attacker now has their session token. From that point on, the attacker can potentially access the victim's entire enterprise account, including email, files and other cloud resources.

Once tokens are stolen, the incident can eventually result in a fraudulent transfer of funds. The "post-compromise" features of the kit include AI tools that help with identifying high-value emails inside the account (such as those involving transfers of funds) and drafting legitimate-sounding responses from within the account. This is reason enough to worry – but the compromise that leads to payment fraud also opens the door to document theft, exposure of regulated data and the ability to reach further into the environment to deploy malware.

Attacks beget attacks

Certain attack types have a viral tendency, especially in today's fragmented "attack-as-a-service" criminal ecosystem. One of the reasons we've seen substantial growth in BEC claims could simply be that threat actors – who may have been working on something else, like exploiting vulnerabilities in VPNs, a few months ago – have noticed the effectiveness of kits like Kali365 and shifted their focus. The proliferation of kits through dark web marketplaces lowers the barrier of entry and expands the number of criminals capable of jumping on any given trend.

But that only tells part of the story. This also serves as an example of AI tools impacting the abilities of attackers in ways that are not at all related to discovering and exploiting vulnerabilities. Phishing kits are dependent, to an extent, on the quality of the lures. The examples of emails researchers have uncovered are highly credible in appearance, with AI likely providing an assist in design. As we mentioned above, the later stages of the BEC attack sequence are directly enabled by AI, and the whole kit is made more attacker-friendly through its intuitive interface, lowering the barrier to entry in terms of coding skills. In this case, AI has not shifted the paradigm – we have had AiTM attacks for years – but nonetheless has had a tangible impact on outcomes, as we see in the rising claim numbers.

Recommendations to defend against token theft

Treat identity, not just email, as the unit of compromise. Assume that access to one enterprise account can expose email, files, collaboration tools and cloud applications together.

Review whether device code authentication is needed. Where it serves no business purpose, restrict or disable the flow and apply conditional access controls to limit exposure.

Monitor beyond passwords and endpoints. Watch for unusual OAuth activity, unexpected token issuance, new application consents and sign-ins from unfamiliar locations.

Build token response into the incident plan. A password reset is not enough — revoke active sessions, invalidate refresh tokens, review OAuth consents and remove unauthorized application access.

Keep verifying high-risk requests out of band. The main goal for attackers is a fraudulent transfer; independent confirmation of payment and banking changes remains a crucial line of defense, whatever technical safeguards are in place.

Briefing: Cyber Litigation Updates – and Three Misconceptions

For several years, litigation over the use of website tracking technologies has been one of the most active and closely watched areas of exposure for businesses.

Years ago, privacy actions mostly centered on customer data exposed by breaches. But as new state laws expanded the rights of website visitors and plaintiff firms tested older privacy statutes in internet cases, we've seen an explosion in actions against companies that had experienced no incident relating to their data.

The basic idea is that pixel trackers, cookies and session replay tools – all commonly used by marketers and website managers at organizations of all sizes – are capturing and transmitting user data without consent. Legal action can be brought against users of these tools under a patchwork of state and federal statutes that may carry significant monetary damages per infringement.

Recently we've seen a step up in plaintiff activity coupled with the seesawing of courts both narrowing the interpretation of some statutes and broadening others – sometimes inconsistently across jurisdictions.

A still-growing litigation wave

The pace of filings is higher than ever in 2026. The Fisher Phillips Digital Wiretapping Litigation Map, which tracks matters filed nationwide since 2022, found 1442 matters filed in the first half of 2026 – 45% more than the same period last year, and four times the number seen in the first half of 2024.

Travelers' own claims data reflects this same upward trajectory. Claims tied to tracking pixels have risen in 2026, driven in significant part by a single serial plaintiff — an individual whose filing activity has since been deemed vexatious by a federal court in California. That plaintiff accounted for much of the recent spike, particularly in Q2 2026, when the number of claims linked to him alone nearly matched the total number of pixel litigation claims Travelers received in the entire preceding year.

But the trend is not the product of one actor. Excluding claims linked to the most prolific plaintiff, claims across the first two quarters of 2026 still ran nearly 60% higher than the quarterly average for 2025. And whereas healthcare entities were once the primary focus of pixel litigation, the exposure is now far more evenly distributed across the economy — healthcare now accounts for just 13% of Travelers' pixel claims, with manufacturing and consulting each representing comparable shares.

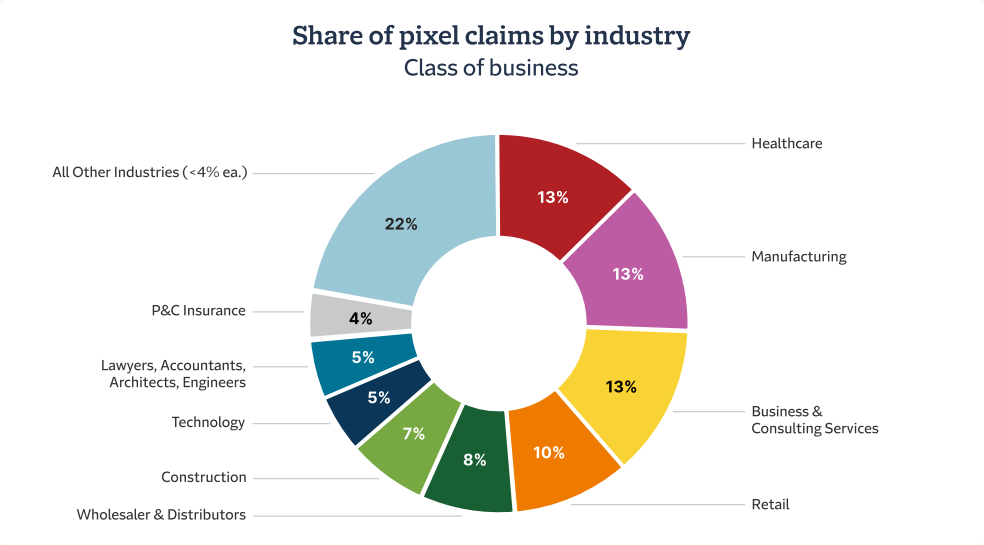

Briefing: Cyber Litigation Updates – and Three Misconceptions (Continued)

Old laws, new applications

Much of this litigation continues to run on statutes written long before the modern web. The California Invasion of Privacy Act (CIPA), enacted in 1967 to address telephone wiretapping, remains the most frequently invoked.

Plaintiffs employ CIPA's wiretap provisions and, lately, its pen-register and trap-and-trace provisions. These provisions originally related to various aspects of wiretapping, from the actual listening or recording of phone conversations to the recording of information about what numbers a person called or received calls from, and the number they themselves used. Plaintiffs argue that routine web tracking amounts to this kind of information capture. (If that seems like a stretch, consider the incentive of $5,000 per violation stipulated by the law.)

Alongside CIPA, plaintiffs have found new uses for other legacy statutes, including the Video Privacy Protection Act of 1988, a federal law now applied to platforms that share users' video-viewing data, and the Illinois Biometric Information Privacy Act (BIPA), which continues to generate class-action exposure for companies handling biometric identifiers.

The state of the Golden State

The outlook is not all bad for defendants, though. The legal foundation beneath these claims — particularly the pen-register theory — is unsettled. Different courts within California have begun to diverge in their readings of CIPA, and the California legislature is weighing a change to the statute itself.

To learn more about these developments and what they mean for organizations, we spoke with Carolyn Purwin Ryan, a Partner at Mullen Coughlin and Chair of the firm's litigation practice. She outlined three misconceptions that she's been hearing in light of recent news around CIPA.

Misconception #1: Businesses not based in California don't need to worry about CIPA.

This remains one of the most dangerous assumptions Carolyn sees. A company does not need to be headquartered in California, have employees there, or maintain any formal California presence to be pulled into CIPA litigation. If the company operates a website that is available nationwide, a single visit from someone located in California may be enough for a plaintiff to try to manufacture a claim. In Carolyn's view, businesses should not treat geography as a shield simply because their physical operations are elsewhere.

At the same time, Carolyn emphasized that courts are beginning to scrutinize these claims more closely, particularly when the plaintiff's only connection to the business is a "drive-by" website visit. Courts in California have increasingly limited standing to plaintiffs who can point to a real account, transaction or relationship with the defendant — not merely a website visit engineered to create litigation. That shift matters. It helps cut off the worst forms of serial-plaintiff activity and gives defendants stronger grounds to challenge cases brought by individuals who suffered no concrete privacy harm.

But Carolyn's bottom line was clear: any organization whose website can be reached from California should assume it may be targeted and should evaluate its tracking practices accordingly.

Briefing: Cyber Litigation Updates – and Three Misconceptions (Continued)

Misconception #2: A cookie banner is enough to shield companies from liability

Carolyn noted that some organizations view having a banner on their website that asks users to accept or reject tracking as a cure-all for CIPA issues. Not so: the banner needs to have the right technical configuration to match the requirements of CIPA. Tracking scripts should not be active until the banner has been interacted with by the user and should honor the user's choice if they reject tracking. A banner pulled "out of the box" in a content management system might not behave properly.

For relatively modest up-front cost, organizations can work with an expert who knows how to configure their cookie banner properly with their tracking tags, Carolyn notes – certainly modest when compared to the potential damages of a CIPA suit. Some organizations opt to do geographic targeting, where a different banner is shown to anyone situated in California, to avoid the more onerous website experience for users in the rest of the country – though Carolyn cautions that geo-targeting isn't perfect and carries some risk that should be weighed by the business leadership.

Misconception #3: New legislation and recent findings will take all of the teeth out of CIPA

As we mentioned under the first misconception, there's a fluid and complex picture around who has standing on CIPA claims, and what threshold for harm needs to be proven for cases to move forward. The general trend is toward narrower interpretations of CIPA, but Carolyn notes that plaintiffs have found new avenues to test before, and likely will again.

Proposed legislation, SB 690, aims to narrow CIPA's application to website activity. Carolyn's view is that SB 690 is important, but businesses should not overread it. If enacted, the bill may meaningfully narrow the pen-register and trap-and-trace theories that plaintiffs have been pressing under CIPA, and it could reduce the number of cases that survive judicial scrutiny. But it will not make CIPA disappear. It will not eliminate the plaintiff bar's incentive to keep testing new theories. And it should not be treated as a reason to pause website-tracking compliance work while waiting for the law to change.

Clearly there are a lot of moving pieces between court findings and legislation – too much for most organizations to keep up with. That's why Carolyn stressed the point that the courthouse is only part of the risk. Much of this activity begins with demand letters, not filed complaints, and those letters can create settlement pressure even when the legal theory is vulnerable. A narrower CIPA may give defendants better arguments, but it will not necessarily stop plaintiffs from sending demands, threatening arbitration or trying to extract nuisance-value settlements. From Carolyn's perspective, the safer assumption is that CIPA will remain a pressure point for website operators—and that companies should prepare for a shifting litigation landscape rather than expect reform to solve the problem for them.

Impact Alerts: Q2 2026

A selection of vulnerabilities gathered by the Travelers Threat Intelligence team.

Travelers issued Threat Alerts to policyholders directly impacted by these vulnerabilities.

On June 8 Check Point disclosed CVE-2026-50751. This is an authentication bypass vulnerability affecting Check Point Remote Access VPN, Mobile Access and Spark Firewall products. It affects deployments configured to use the deprecated IKEv1 key exchange protocol where gateways accept legacy Remote Access clients and don't require a machine certificate for connections. It was added to CISA's KEV Catalog the same day with reports of first exploitation observed May 7, 2026. At least one incident is attributed to a Qilin ransomware group affiliate.

Oracle disclosed a vulnerability, CVE-2026-35273, in PeopleSoft. The vulnerability exists within the Updates Environment Management component of PeopleSoft Enterprise PeopleTools and can be exploited remotely over HTTP without any credentials. The underlying flaw has been classified as a server-side request forgery (SSRF, CWE-918), with successful exploitation enabling remote code execution. This was added to CISA's KEV Catalog on June 12, 2026. Exploitation, attributed to ShinyHunters extortion group, was seen starting May 27, 2026, two weeks prior to Oracle's disclosure. Higher education institutions were highly targeted during this campaign.

In June 2026, a security researcher surfaced a dataset being called "FortiBleed." This was a collection of what appeared to be working login credentials for roughly 74,000 FortiGate firewall URLs spread across 194 countries and over 21,000 domains. Fortinet believes that this isn't a new flaw in their products but instead, threat actors are recycling credentials from two older, previously disclosed vulnerabilities, and pairing that with brute-force password attacks aimed at devices that have weak passwords and no MFA enabled. CISA released an alert on urging the hardening of Fortinet devices. There were reports of data being sold in relation to this leak.

Conclusion

The through-line of the quarter is compression — of time, of cost and of the margin for error. Ransomware volume held near record levels while the field of operators kept fragmenting, making any one group's disruption less consequential to the whole. Account takeovers surged as attackers found a faster path around authentication controls. And the early results from Project Glasswing confirm that finding software vulnerabilities is becoming cheap, shifting the burden of security onto how quickly an organization can verify, prioritize and remediate what is found. Our own claims data still shows no sign that attackers have gained meaningful new capability in vulnerability discovery — for now, Mythos-class capability sits with defenders. But the direction of travel is clear, and the window between a flaw becoming known and becoming exploited is the one thing this technology is certain to shorten.

For organizations, the response is less about any single new defense than about executing familiar ones with greater urgency. Treat patch timelines as hard operational deadlines. Treat identity, not just email, as the unit of compromise. And recognize that independent, out-of-band verification of high-risk requests remains a crucial safeguard whatever technical controls are in place. The threats are evolving, but the fundamentals of resilience have not — and defenders still hold the upper hand where they choose to act on it.

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.