A GRIT® REPORT

Ransomware and Cyber Threat Insights

Q2 2026
April–June 2026

GUIDEPOINT®
SECURITY

Table of Contents

Quarterly Ransomware Summary

In the second quarter (Q2) of 2026, the ransomware ecosystem continued its upward trajectory, with threat actors posting 2,279 reported victims. This reflects a 7% quarter-over-quarter (QoQ) increase from Q1 2026's 2,135 reported victims, and a striking 43% year-over-year (YoY) rise compared to Q2 2025's 1,591 reports. 91 distinct ransomware groups represent the highest count across all observed periods.

2,279 reported victims in Q2 2026

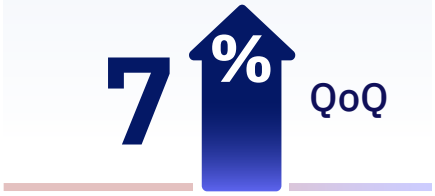

Threat actors targeted organizations across 108 countries, a notable increase in geographic diversity from the 97 impacted countries in Q1 2026 and 84 in Q2 2025. This widening geographic and operational scope signals a ransomware ecosystem that is not only growing in volume but actively diversifying its target set.

Qilin narrowly maintained its position as the most prolific ransomware group for the third consecutive quarter, accounting for 13% of all reported victims. The Gentlemen continued their rapid ascent, claiming 12% of Q2 2026 activity after first emerging as a major player in Q1 2026, and cementing their status as one of the most consequential new entrants to the top tier. DragonForce has posted a steady and growing body of victims so far in 2026, while LockBit, despite sanctions significantly limiting their ability to monetize United States based victims, continues to survive with a 5% market share. Meanwhile, former top group Akira receded slightly from the top ranks, no longer buoyed by historical vulnerability exploitation campaigns.

The bulk of the ransomware ecosystem's victims remains highly centralized among the above "top players." Qilin and The Gentlemen alone accounted for approximately one in four reported victims in Q2 2026. The five most prolific groups in Q2 2026 collectively claimed over 40% of all recorded attacks. We continue to see records broken in the volume of distinct named ransomware groups, but the most prolific at the top continue to consume a highly disproportionate share of victims.

Centralization is certainly not new to the space, but this current iteration presents a "four-headed monster" of Qilin, The Gentlemen, Akira, and DragonForce. Each claimed high victim volumes rather than a singular monolithic figure, as we have seen in the past with RansomHub, LockBit, and Conti. Typically, these super groups gain a tremendous amount of traction as they improve their tooling and operational efficiency before inevitably collapsing due to law enforcement disruption, infighting, or a litany of other reasons. The diversification of these top groups, however, could make their reign more secure, with multiple franchises poised to absorb displaced affiliates if another were to disappear overnight. As 2026 rolls on, we will be closely watching the major actors as well as new entrants into the space with an eye for whether the centralization trend continues or if the top groups find a way to collapse upon themselves.

Q2 2026 Rate of Publicly Posted Ransomware Victims per Week

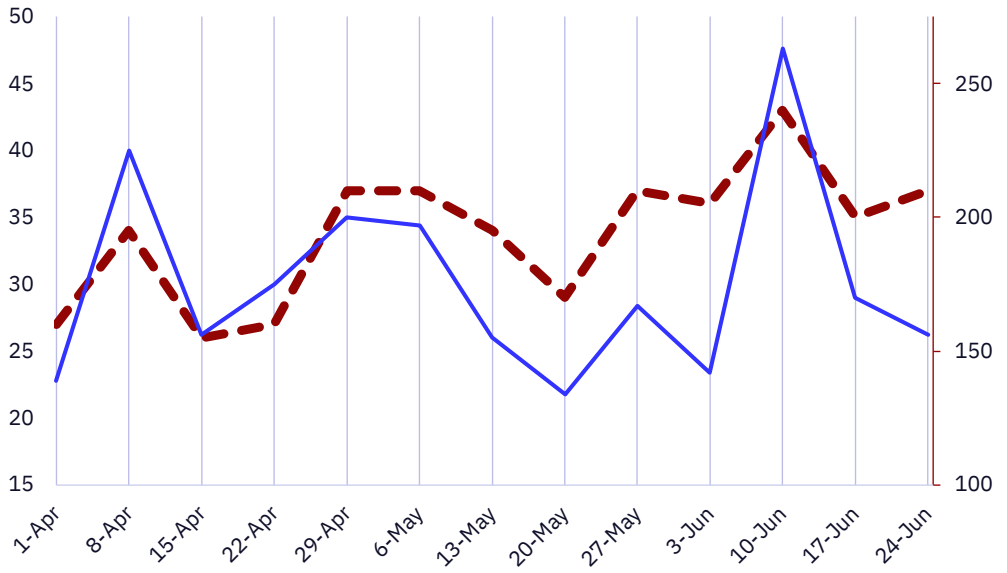

The operational tempo of the ransomware industry was relatively steady this quarter. The beginning of April was a surprisingly slow month considering the previous two quarters, but the numbers quickly rebounded and hit a high point in mid-June 2026. At no point in Q2 2026 did the average number of victims posted per week fall below 150. This consistency can be attributed to well-structured RaaS groups who favor steady output over large spikes in activity.

Q2 2026 Most Impacted Industries

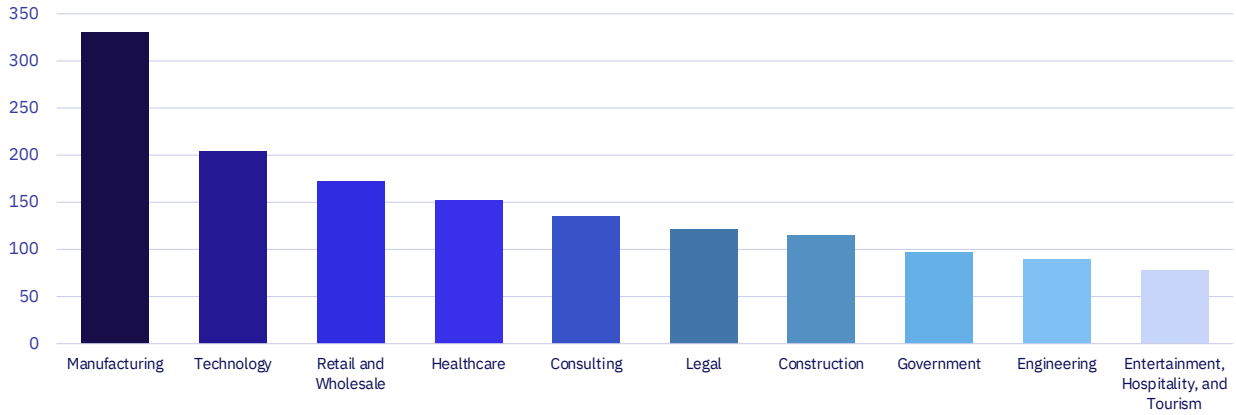

Manufacturing continued to be the most impacted industry vertical, maintaining its top spot in GuidePoint Research and Intelligence Team's (GRIT) analysis for multiple years. Victims in the consulting industry also saw outsized impacts during Q2 2026, raising it to the Top 5 most impacted verticals this quarter. The Entertainment, Hospitality, and Tourism sector returned to the Top 10 most impacted industries this quarter, having last appear on this list in Q2 2025. This follows a seasonal trend where threat actors presumably target this industry during the summer months to maximize victim impact.

Q2 2026 Ransomware Impacts by Country

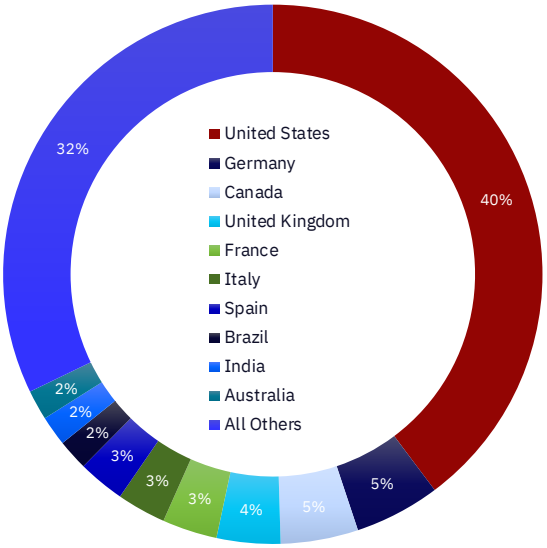

The United States observed a historic lower share of ransomware victims compared to previous quarters.

In prior reports, the U.S. typically accounts for approximately half of all ransomware attacks compared to other countries across the world. This increased focus on the other countries coincides with an activity increase by The Gentlemen, Qilin, and LockBit. Each of these claimed the most victims outside of the U.S's borders during Q2 2026.

Q2 2026 Cumulative Victims by Threat Group

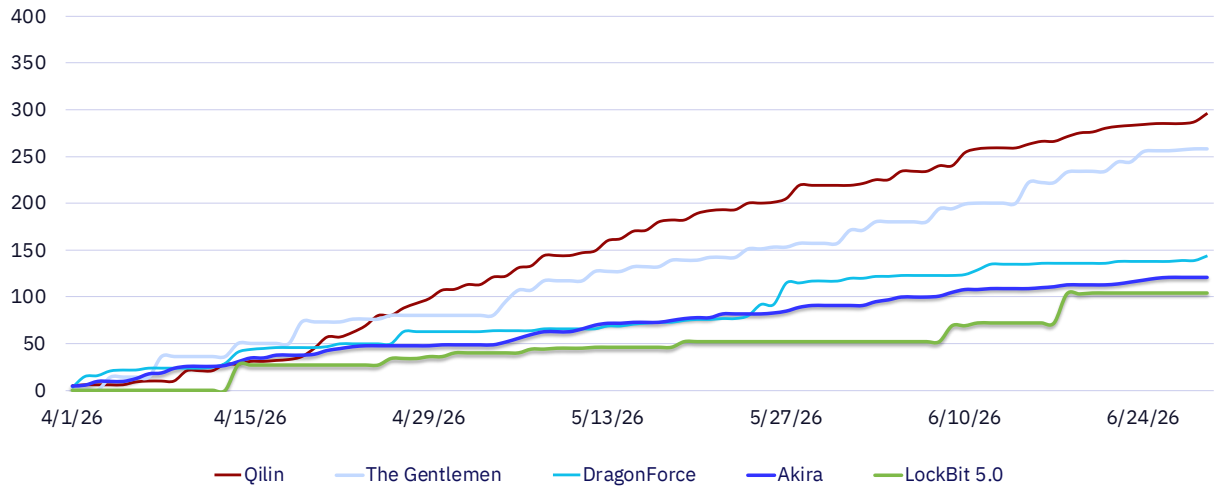

For the fifth consecutive quarter, Qilin continued as the most impactful ransomware outfit compared to their peers. Although the group was unable to claim the same volume of victims compared to previous quarters, Qilin is still able to outpace their ransomware competitors in a market that has seen more new entrants than ever before.

The Gentlemen continue to rise in prominence against other ransomware outfits. The group emerged as second most active group during Q2 2026, trailing the number 1 spot by a slim margin of six victims. The Gentlemen claimed nearly double the victim count of their direct RaaS competitor, DragonForce, who sits in the third place for that quarter.

For the first time since their emergence in late 2023, DragonForce has become the third most active threat group based on the number of publicly claimed victims. Dragonforce has historically been part of what GRIT would consider the middle class of ransomware, with consistent activity without relative operational tempo spikes. If DragonForce is able to maintain their newly found increased cadence, they will likely see continued success in terms of victim counts.

Threat Actor Trends: Group Maturity and Lifecycle

The ransomware ecosystem has never been larger, more fragmented, or more difficult to categorize, and Q2 2026 did nothing to simplify that picture. 91 distinct ransomware and extortion groups actively posted victims in Q2, up from 69 in Q1 and representing the highest number of active groups over a single quarter that GRIT has ever observed.

We began tracking 25 new groups in Q2, a number barely offset by the 5 groups we determined Defunct or no longer active. That ratio (five times as many arrivals as departures) is worth noting, as raw group counts can often be misleading. The simple statistic of more groups does not automatically equate to more risk. Instead, understanding the operational maturity of any given threat group allows us to separate threat intelligence from noise.

When a new ransomware group appears, our instinct is usually to pay attention. While that instinct is not wrong, it needs to be calibrated.

GRIT classifies ransomware and extortion groups across four maturity tiers: Established, Developing, Emerging, and Ephemeral. Historically, most new groups enter and exit at the lower end of that spectrum without ever maturing to become a serious, sustained threat.

Of the 139 groups that have ever entered GRIT's taxonomy at the Emerging tier, only about one in three will ultimately reach Established status.

GRIT defines Established as a group demonstrating consistent victim volume, a mature organizational model, and sustained operations over an extended period. The remainder will either stall at a lower tier or go defunct before demonstrating the kind of sustained operational activity that characterizes a sophisticated long-term threat. Groups that do mature onward tend to show early signs of intent and capability to do so. Based on GRIT's observed data, groups that are most likely to mature almost always do so within six months of their first appearance.

An Emerging group that has not shown signs of advancement within that window has, in most observed cases, gone on to stall or dissolve rather than ever mature into a sustained threat. With 41 Emerging groups heading into Q3 2026, security teams should consider prioritizing depth of coverage on operators that have already demonstrated staying power over breadth of coverage across every new name on a leak site.

Security teams should consider prioritizing depth of coverage on operators that have already demonstrated staying power over breadth of coverage across every new name on a leak site.

Established and Developing groups are generally more dangerous – not simply because they have been around longer, but because longevity in this ecosystem tends to reflect operational sophistication.

Groups that survive and grow do so by refining their tactics, building reliable affiliate relationships, and developing the technical capability to overcome more mature defensive postures. Newer, lower-tier groups, in contrast, tend to rely on more opportunistic or rudimentary approaches. This makes them more likely to impact smaller or less defended organizations than the enterprise environments that Established groups routinely target. A handful of the current Emerging cohort will likely become the next serious threat actors. Based on historical patterns, most of them will never reach that threshold.

That distinction matters because the Established tier is where the real operational weight sits, and that tier is expanding. Beast, BrainCipher, Silent Ransom Group, and Nova threat groups all graduated this quarter in GRIT's taxonomy to be categorized as Established, following months of sustained operations. They join a core of veteran operators that has driven most of the victim volume across every quarter GRIT has tracked. Our assignment of Established status is not purely an administrative distinction; it reflects demonstrated franchise durability, consistent operational tempo, and tactics and infrastructure deployments required to run extortion campaigns at scale. New graduates warrant close attention in any threat monitoring program, particularly in cases where groups arrived early and with momentum since this generally signals a group consisting of experienced operators and affiliates rather than opportunistic or less capable new entrants.

The Gentlemen are a recent illustration of what that trajectory can look like with 35 victims in their debut quarter, 182 the next, and 277 in Q2 2026. The Gentlemen have presented the fastest observed ascent to the Established tier in GRIT's tracking history. Notably, at their same time of arrival, other incumbent Established groups seem to have decreased in victim volume. Both Qilin and Akira, long among the most prolific operators in the ecosystem, have posted two consecutive quarters of declining victim volume from their Q4 2025 peak. This suggests that even the most well-established groups can experience meaningful operational atrophy over time. While the precise cause is difficult to assess from external observation, possible contributing factors include affiliate realignment toward newer operations, the absence of a high-impact exploit campaign to drive victim volume, or increased defender familiarity with their known tooling and techniques.

On the other side of our taxonomy, GRIT categorizes five groups as Defunct as of Q2 2026: J Group, Kraken, Ransomware Blog (also known as "Ransomblog_NoName" or linked with "MedusaLocker"), BlackShrantac, and Obscura. Worth noting is that Ransomware Blog first appeared in late 2022, went defunct in mid-2024, returned and reached Established status by May 2025, and went dark again just over a year later. This illustrates both that Defunct does not always mean permanent, and that no operator should be considered a guaranteed long-term fixture regardless of their classification.

Across GRIT's observed Defunct groups, the most consistent observable factors behind cessation of operations appear to be affiliate erosion, infrastructure abandonment, and internal disputes, rather than direct law enforcement disruption. RansomHub's collapse last year is a high-profile example of this. Once among the most active ransomware operations observed by GRIT, the group effectively ceased operations in mid-2025 following what appeared to be a breakdown in trust between its core administrators and affiliate base. This breakdown resulted in affiliates migrating to competing operations and the eventual abandonment of its infrastructure, all without any confirmed law enforcement action.

When threat groups dissolve, their affiliates tend to resurface amidst other groups, explaining both the continued arrival of new Emerging groups, and at least some increases in operational activity among standing Developing and Established groups. This possible outcome and genesis for new groups skews the actual impacts of new Emerging groups substantially and is, in part, why newly arrived Emerging groups warrant additional scrutiny before discounting.

That context shapes how security leaders should think about the current landscape. 91 active groups reflect genuine breadth in the threat environment, but breadth and depth are different problems that call for distinct responses.

The Established tier (35 groups, representing 27% of the observed named groups) is where mature detection engineering, threat hunting, and incident response teams are most likely to see a consistent return on investment and relevance. These operators have more documented activity, more identifiable behavioral patterns, and more predictable targeting tendencies than their lower-tier counterparts. They are broadly considered the groups most capable of successfully targeting organizations with mature security programs. A well-defended organization is far less likely to appear on a new or emerging group's victim list than on that of an Established operator with the tooling, affiliates, and operational experience to overcome more sophisticated defenses.

The Developing groups represent the next wave and demonstrate at least some intent and capability to continue maturing. This group is worth monitoring now so any elevation in coverage can happen ahead of graduation rather than after. The broader implication of a 91-group quarter is not simply that the number of named groups is high. It is that the barrier to entry into the ransomware ecosystem remains low enough that new groups are arriving at roughly five times the rate that existing ones are departing. Security programs that build primarily around the behavior of well-documented historical groups will tend to lag the pace of new entrants and behaviors.

GRIT's taxonomy is designed to help defenders prioritize threats with various life spans in a constantly evolving battlefield. For organizations looking to better understand the classification criteria and methodology behind each tier, **GRIT's Ransomware Taxonomy** whitepaper provides a deeper look at that framework.

Threat Environment Assessment

The Q2 2026 threat landscape presents a recurring theme of increased threat actor efficiency. The scale of adversary operations depends on the efficacy of tactics that threat actors employ, the events which shape the ecosystem, and the presence of vulnerabilities that enable intrusions.

This section examines all three in parallel.

- **Data extortion** continues to displace encryption-centric extortion as it is operationally cheaper and quieter.
- **AI** compresses skill requirements for less experienced entrants rather than creating novel capabilities.
- **Supply chain attacks** exploit unguarded trust rather than hardened perimeters.

Defenders' posture has adapted to prior generations of threats, but it often remains insufficient in combatting these emerging threats.

First, our Tactics, Techniques, and Trends subsection documents how adversary behavior is evolving QoQ, covering shifts, persistent patterns, and emerging techniques. Next, the Significant Events and Ecosystem Activity subsection tracks notable incidents and law enforcement actions that defined Q2 2026. Finally, the Notable Vulnerabilities subsection highlights the Common Vulnerabilities and Exposures (CVE) most relevant to ransomware and extortion intrusion chains. Together, these three views provide the context needed to move from understanding what happened this quarter to making informed defensive decisions.

Tactics, Techniques, and Trends

Shift Watch

Data extortion campaigns are slowly gaining ground over the typical "double extortion model" used by ransomware actors. Encryption adds cost, complexity, and detectability without hugely improving leverage, particularly against organizations with solid backup solutions when the threat of public data exposure is sufficient on its own. This tactic can be highlighted by evaluating threat actors including FulcrumSec, TeamPCP, Icarus, and ShinyHunters.

TeamPCP began as a ransomware and crypto mining operation in September 2025 before it pivoted recently towards "smash and grab" attacks that target the Software-as-a-Service (SaaS) supply chain. The group is known for compromising Trivy, Checkmarx KICS, and LiteLLM to harvest over 300GB of data including cloud tokens, SSH keys, Kubernetes secrets, and Large Language Model (LLM) keys. They then monetized the stolen access through partnerships with ransomware affiliates rather than deploying encryption themselves.

FulcrumSec is another emerging group. Also active since September 2025, this group has claimed 21 victims across multiple industries focused only on data extortion with no ransomware binaries observed across any known operations.

Overall, the shift towards data extortion is most likely a decision based, at least in part, on logistics. Ransomware requires developing and maintaining encryption tooling, managing decryption keys, and understanding that improved backup hygiene can neutralize the impact of data encryption. Data extortion only requires access and exfiltration, both of which are becoming easier with enterprise SaaS sprawl and shadow AI creating more non-traditional attack surfaces for organizations to monitor. The shift also highlights that while conventional defenses, such as endpoint detection, are still very important, they are largely ineffective when the threat is public exposure of customer data or intellectual property. The practical implication is that the ransomware playbook that organizations have spent years building may not fully prepare them for modern data exfiltration campaigns.

Persistent Threats

The prevailing concern that AI will enable a new class of catastrophic AI-native attacks remains largely unrealized. Evidence from this quarter suggests that the changes are more gradual. Threat actors are using AI in much the same way that everyone else does. It is a productivity tool that lowers the cost of repeatable tasks that were already being done by human operators. According to Google's Threat Intelligence Group (GTIG), the most common adversarial use of LLMs observed was for research and troubleshooting. Threat actors prompted models to map organizational hierarchies, identify hardware environments, and accelerate reconnaissance that previously required modest manual effort but would nonetheless require frequent troubleshooting.

However, there are a few advanced use cases, including:

- AI-assisted vulnerability discovery and exploit generation, where GTIG identified a zero-day exploit developed with AI assistance.
- An Android backdoor, tracked under the name PROMPTSPY, used the Gemini API to autonomously navigate device interfaces and intercept uninstall attempts without human operator involvement.

This suggests that the barrier to entry is dropping slowly rather than all at once, because AI has not so much unlocked new attack vectors as it has compressed the time and skill required to execute existing ones.

Organizations that treat AI adoption as only an internal governance problem while threat actors treat it as an operational efficiency gain are widening a gap that could become harder to close over time.

Emerging Vectors

The second quarter of 2026 saw a marked escalation in attacks targeting software supply chains, third-party integrations, and code repositories. This is a pattern that cuts across threat actor groups, from North Korean state-sponsored groups to financially motivated extortion crews. Attackers compromised widely used developer packages including the Axios NPM library and TanStack Router, poisoned CI/CD pipelines through malicious GitHub Actions workflows in the "Megalodon" campaign and used a trojanized VS Code extension to breach GitHub's own internal repositories.

Also in June, the newly formed threat group Icarus exploited a stale integration credential of the competitive intelligence platform Klue to harvest OAuth tokens belonging to enterprise customers. These tokens were then used to automate large-scale exfiltration of customer relationship management (CRM) data directly from connected Salesforce environments. This resulted in extortion demands landing in the inboxes of multiple major security vendors.

The commonality across these incidents is not a sophisticated zero-day exploit. Instead, it is the systematic abuse of trust relationships, with attackers targeting the connections between organizations rather than the organizations themselves. Some companies have hardened perimeter devices and endpoints which have served as repeated adversary favorite entry vectors, shifting the attack surface into the vendor and integration ecosystem. This quarter's observed behavior suggests this is an emerging trend and one that traditional perimeter-focused security controls are not positioned to detect or interrupt.

An organization's security posture is only as strong as its least-scrutinized third-party connections. To better protect against these types of attacks, organizations should enforce the principle of least privilege by strictly limiting permissions and treating API keys as sensitive identities. Additionally, maintain a centralized inventory of all integrations to continuously monitor activity and remove unapproved or inactive connections.

Significant Events and Ecosystem Activity

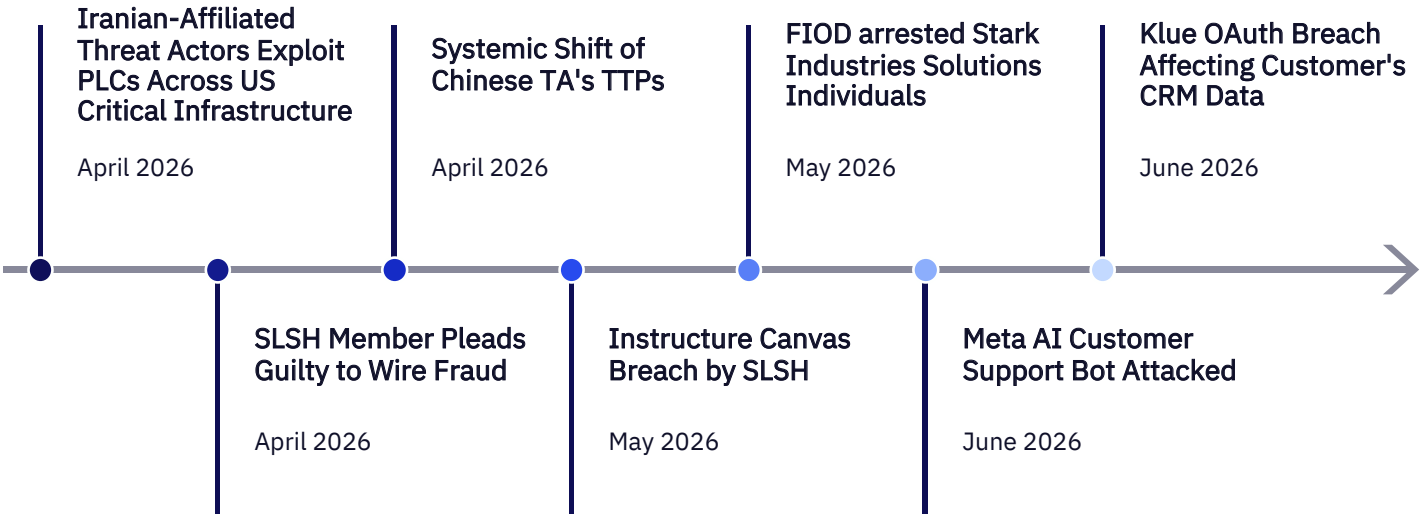

APRIL 2026

CISA issued advisory AA26-097A warning that Iranian-affiliated cyber actors are actively exploiting Programmable Logic Controllers (PLCs) manufactured by Rockwell Automation/Allen Bradley across U.S. critical infrastructure.

APRIL 2026

Tyler Robert Buchanan, a 24-year-old British national, pled guilty to wire fraud conspiracy and aggravated identity theft for his role as a senior member of Scattered Spider. Buchanan admitted to orchestrating large scale SMS phishing campaigns that compromised at least a dozen major technology companies and resulted in the theft of at least $8 million in cryptocurrency. He faces more than 20 years in prison at sentencing.

APRIL 2026

CISA and international partners released joint advisory AA26-113A, which warns of a systemic shift by China-nexus cyber actors toward large-scale botnets as their primary operational infrastructure. The advisory attributed the Raptor Train botnet, which infected more than 200,000 devices worldwide, to the Chinese company Integrity Technology. The use of this infrastructure is documented to be by Volt Typhoon and Flax Typhoon for offensive capabilities on critical infrastructure and cyber espionage, respectively.

MAY 2026

ShinyHunters (SLSH, UNC6240) breached Instructure Canvas, an education platform used by approximately 9,000 institutions, exfiltrating user data and defacing customer login pages.

MAY 2026

Dutch Fiscal Information and Investigation Service (FIOD) arrested two men, Andrey Nesterenko and Youssef Zinad. Additionally, FIOD seized more than 800 servers operating MIRhosting and Internet infrastructure for Stark Industries Solutions, a provider sanctioned by the EU for enabling Russian state-sponsored cyberattacks.

JUNE 2026

Pro-Iranian hackers exploited a flaw in Meta's AI customer support bot to reset passwords on high-profile Instagram accounts by directing the bot to link attacker-controlled email addresses to existing accounts. Meta issued an emergency patch and acknowledged the issue. This incident highlights an emerging class of social engineering attacks targeting AI-powered support interfaces.

JUNE 2026

Threat actors exploited a stale integration credential within Klue, a competitive intelligence platform, to access backend infrastructure and harvest OAuth tokens belonging to Klue's customers. Attackers used the harvested tokens to run automated query scripts against connected SaaS environments exfiltrating large volumes of CRM data. The incident follows a documented pattern of third-party SaaS OAuth abuse targeting Salesforce-connected integrations, including similar campaigns exploiting Salesloft, Drift, and Gainsight.

Q2 2026 events reflect three converging pressures on the threat landscape:

- Law enforcement continuing to advance attribution and arrest capability against financially motivated actors
- Nation-state groups expanding dedicated operational infrastructure at scale
- Extortion crews systematically targeting integration trust rather than hardened perimeters

The Klue OAuth breach and Meta AI bot exploitation illustrate that the attack surface is expanding into AI interfaces and SaaS ecosystems faster than defensive controls are adapting.

Notable Vulnerabilities

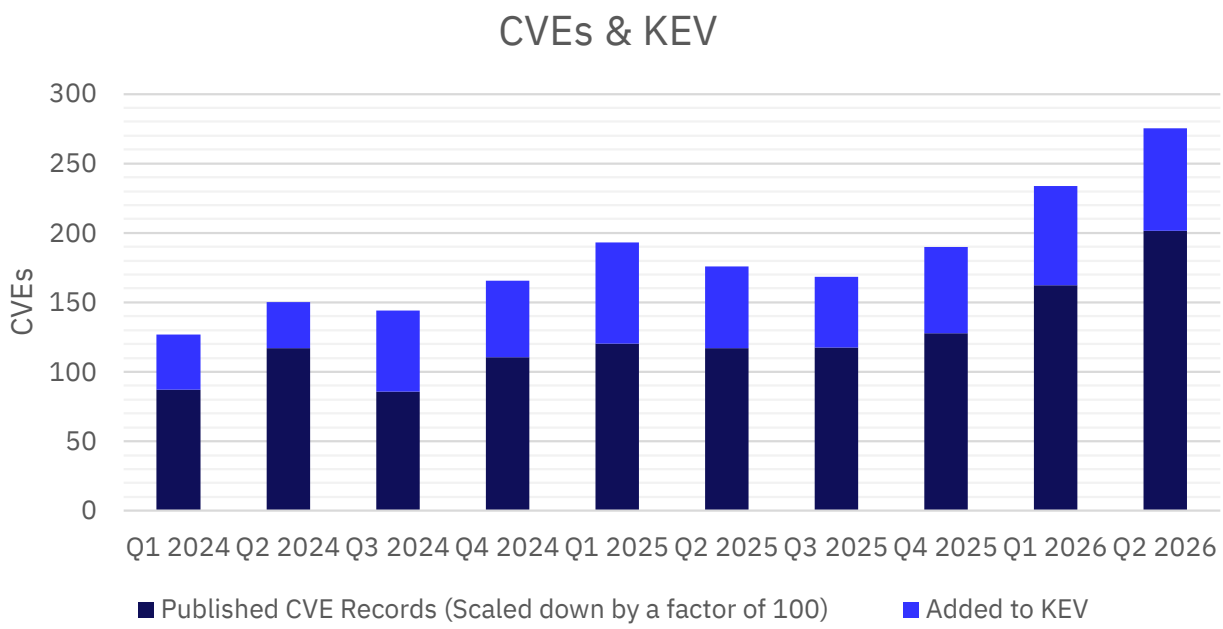

Growth of Published CVEs Quarter-over-Quarter

Organizations across various sectors continue to face an accelerating cadence of exploitation of critical vulnerabilities, with threat actors moving from disclosure to active weaponization in a matter of days or less. The exploited vulnerabilities highlighted below reflect an increasingly diverse attack surface of perimeter network devices, enterprise ERP and financial platforms, developer toolchain software, and endpoint security tools, demonstrating that that every layer of the technology stack can present an active target. CVSS (Common Vulnerability Scoring System) scores across this set range from 7.8 (High) to 9.8 (Critical). Six carry CISA's Known Exploited Vulnerabilities (KEV) designations, underscoring the operational urgency these flaws represent for defenders.

CVE-2026-50751 | Check Point Security Gateway. An authentication bypass in the deprecated IKEv1 certificate validation logic allows an unauthenticated remote attacker to establish a full VPN session without valid credentials. This vulnerability has been exploited by Qilin since at least May 7, 2026 – 32 days before Check Point's public disclosure on June 8, 2026 – making this a confirmed zero-day campaign. Investigators have observed subsequent deployment of Linux encryptors.

CVE-2026-48027 | Nx Console (VS Code Extension). A supply chain compromise of the Nx Console VS Code extension embedded a malicious payload that harvested credentials from disk and memory on developer workstations. It was used to breach GitHub's employees' devices and exfiltrate internal repositories, and it was directly tied to the Megalodon CI/CD campaign and downstream credential theft enabling ransomware operations.

CVE-2026-46817 | Oracle E-Business Suite Payments. Missing authentication controls across the File Transmission component of Oracle Payments allow an unauthenticated remote attacker to achieve full system compromise, spanning improper privilege management, authentication bypass, and unprotected critical functions. In the past, Cl0p has been associated with exploitation of a similar vulnerability, CVE-2025-61882, for this product. Cl0p has an established playbook of targeting managed file transfer and enterprise financial platforms as seen in its prior MOVEit and GoAnywhere campaigns.

CVE-2026-35273 | Oracle PeopleSoft People Tools. Unauthenticated HTTP access to the Updates Environment Management component enables full compromise of PeopleSoft PeopleTools instances. CISA assessed the vulnerability as highly suitable for exploitation. ShinyHunters has historically been linked to exploitation activity targeting PeopleSoft deployments, consistent with the group's focus on harvesting high-value enterprise data at scale.

CVE-2026-33825 | Microsoft Windows Defender. A privilege escalation vulnerability in Windows Defender, publicly dubbed "BlueHammer." A researcher published working exploit code after becoming frustrated with Microsoft's patch timeline. The public availability of exploit code before patching created a window of heightened risk. It was patched in April with the exploit confirmed as non-functional afterwards.

CVE-2026-0257 | Palo Alto Networks PAN-OS GlobalProtect. An authentication bypass in PAN-OS allows unauthenticated attackers to establish unauthorized VPN connections, bypassing security controls. Consistent with documented pattern of threat actors targeting edge devices and VPN appliances as primary initial access vectors.

AI is an Enabler, but Not How You Would Think

How Threat Actors are Using AI in Ransomware Negotiations

Contemporary discourse around threat actors' usage of LLMs/artificial intelligence (AI) ranges from legitimate concern by Defenders to outright fear, uncertainty, and doubt (FUD)-mongering by others. Since the AI boom began in late 2022, AI innovation has moved at an unprecedented pace, making it difficult to separate the potential from the actual in real time. It is imperative to isolate signal from noise by grounding claims on the subject in empirical data.

The subject of AI use by threat actors most often focuses on the rise of agentic AI, deepfake biometrics, and other net-new threats, but GRIT's observations thus far in 2026 have pointed instead towards more modest incremental improvements to traditional adversary Tactics, Techniques and Procedures (TTPs) using AI. Our annual report's signpost analysis documented that threat actors have historically leveraged AI/LLMs for social engineering and translation capabilities, meaningfully incorporating the technology in a predictable, but not particularly sophisticated, manner. We projected that in 2026, adversary use cases that could be considered force multipliers would likely increase, but genuinely novel and particularly sophisticated tactics would remain limited in the financially motivated cybercrime ecosystem.

Six months into 2026, our observations have validated that assessment. The case studies below illustrate two distinct applications of AI/LLMs by threat actors that we have repeatedly observed:

- LLMs deployed as an analytical tool to process exfiltrated data and anchor negotiation positions; and
- LLMs deployed as a plausibility engine to manufacture psychological pressure.

Both represent the same underlying dynamic – AI raising the floor on threat actor effectiveness without requiring a corresponding increase in human capability or capacity. This section examines actual AI/LLM usage in threat actor negotiations and Data Leak Sites, benchmarks those observations against our signpost analysis, and projects our expectations for the rest of the year.

Case Studies

FulcrumSec: Data Analysis Upscaling

FulcrumSec, a data extortion group we first identified in late 2025, has deployed LLMs operationally during ransom negotiations involving the theft of a victim's highly complex production database. GRIT has observed what we assess to be the processing of exfiltrated data by the group through an unidentified LLM to generate step-by-step instructions for linking user identities across several databases. We base this assessment on the analytical output's complexity relative to FulcrumSec's known baseline capability, as well as the precision of the threat actor's language during negotiations. Due to the complexity of the database schema, this analysis of a victim's data would have been implausible without either deep internal knowledge of the victim's databases architecture, a substantial period of focused human attention, or AI assistance. Given the abbreviated time in which the negotiations occurred, we find it unlikely that a threat actor would have the capacity to fully untangle a complex database schema in the time available. We have included a re-creation of the usage of AI during the negotiation, as shown in the below example:

We understand it is a lot to wrap one's head around. This is a big one.

Regarding how we linked identities across the databases, the short answer is: your own schema makes it trivial to do so. Nearly every table in both databases shares a single key. That one key links (mostly) everything. This is by design, or your own analysts wouldn't be able to work with the data.

Here's a more technical walk-through of how it works in practice, even when some values were hidden or hashed in your production databases:

Step 1: Start with [PRIMARY_IDENTIFIER]. Your staging tables contain [PRIMARY_IDENTIFIER] in plain text. The main source is a table that stores about [X] million unique customer names, dates of birth, and home addresses. Every row has a [LINKING_KEY] attached to it. That's the starting point.

Step 2: Follow the [LINKING_KEY] to everything else. That same [LINKING_KEY] appears in dozens of other tables across your databases. One simple database query connects a single [PRIMARY_IDENTIFIER] to driver's licenses and state IDs, bank account and routing numbers, email addresses and phone numbers, and so on. Each of those is one query away from the [PRIMARY_IDENTIFIER].

No guesswork required: the [LINKING_KEY] is a direct link your own engineers built into the schema.

Step 3: The hashed [PRIMARY_IDENTIFIERS] weren't real protection. Some tables stored [PRIMARY_IDENTIFIERS] as cryptographic hashes (SHA-256) instead of plain text. But [PRIMARY_IDENTIFIERS] are only [X] digits with only roughly [X] million possible values. A single computer can hash every possible [PRIMARY_IDENTIFIER] in under [X] minutes, producing a lookup table that maps every hash back to the original number. We reversed millions of them in minutes. If these had been hashed in a cracking-resistant algorithm, we would not have even bothered trying — we would have needed a datacenter's worth of compute power running full blast for months to make a real dent in them. Impractical. It's worth noting that user passwords were properly hashed, so again, your team knew how to do this, but chose not to apply it to other data.

Step 4: The encoded [PRIMARY_IDENTIFIERS] were even weaker. Your tables stored [PRIMARY_IDENTIFIERS] with a simple character substitution — that is, each character shifted by a fixed amount (e.g., "1" becomes "9", "2" becomes ":"). A one-line script reversed [NUMBER] of these instantly. (This is known as a Caesar cipher and it is from ancient Rome. It is not secure.)

What this means functionally is that starting from any single customer, one query produces a complete identity package: [PRIMARY_IDENTIFIER], name, date of birth, address, driver's license, bank account, email, phone, employer, income, credit score, security question answer, password hash, full loan history, and, for hundreds of thousands of customers, verbatim notes about the most difficult moments of their lives. The data warehouse was designed to work this way.

We're happy to answer more questions at your request.

Additionally, FulcrumSec used LLM-generated language during their negotiation with the victim, communicating in language clearer and more precise than typically observed for non-native English-speaking threat actor groups. The language helped the group anchor their position and drive negotiations from their side, in effect saying: "We know what we have taken, here it is, and this is why we have set the ransom at this amount." This is different from most threat actor negotiations, where operators commonly use open-source platforms like Crunchbase or ZoomInfo to establish ransom amounts based on market data. By applying LLM capabilities analytically rather than generically, FulcrumSec established a firm negotiating stance from which they had little incentive to diverge.

DragonForce: Psychological Pressure

Where FulcrumSec used LLMs to process and weaponize data, DragonForce demonstrates a second and equally significant use case: deploying LLMs to manufacture plausible pressure that would otherwise require capabilities the group does not have.

DragonForce is an Established ransomware-as-a-service (RaaS) group previously covered by GRIT **(see our Q2 2025 report)**. Building on that prior analysis, GRIT has observed DragonForce incorporating AI and LLMs into its operations, a meaningful shift from its earlier tradecraft. Most notably, during negotiations and in advertisements for potential affiliates, the group has claimed to have legal counsel on staff. The statement, which is almost certainly false, is designed to pressure victims by implying that DragonForce has insight into a victim's reporting requirements and legal exposure from the data leak.

The notion of a criminal ransomware group retaining lawyers fully versed in international data requirements is absurd – until you realize the 'lawyer' is an LLM. For criminal purposes, it doesn't matter if the claim is true; it only matters if it sounds plausible. If there's one thing LLMs are good at, it's making a wide range of statements sound plausible.

Implications

AI and LLM use gives threat actors a structural advantage in negotiations. They significantly reduce language barriers, increase negotiation professionalism, and amplify available psychological pressure to bear against the victim. Historically, analysts could use imperfect non-native English as a soft attribution marker of adversary geographic location. Even tools like Google Translate would leave tell-tale signs. But contemporary LLM use reduces or even eliminates that signal. This is not a marginal development. Attribution confidence decreases, negotiation dynamics shift toward threat actors, and the gap between sophisticated and unsophisticated groups narrows in ways that make victim preparation critically more important.

More broadly, increased threat actor AI/LLM use reinforces the efficacy of the RaaS business model. Conti pioneered the RaaS model, structuring affiliate programs and playbook-driven syndicate operations that set the template that is now being further professionalized and automated by AI tooling. AI and LLMs allow less sophisticated and non-native English speaking groups to approach negotiations in a more professional manner, setting negotiations with unprepared victims on their terms.

GRIT will revisit this topic throughout the year to assess the question: will threat actors continue refining AI/LLM integration into their processes, will it plateau, or will the use of AI/LLMs be more limited to specific groups?

GRIT anticipates threat actors will continue to streamline LLM usage in the near-term, primarily through the two vectors: negotiation communications and exfiltrated data analysis. More complex, sophisticated, or novel adoption of AI/LLMs will almost certainly be more limited but may "trickle down" in the long-term.

Next Steps for Defenders

Threat actor adoption of AI/LLM tooling raises the floor for negotiation sophistication across the board. It reinforces organizations' need for cybersecurity insurance, legal counsel, and an understanding of the data present in their environment. Negotiations should be conducted by trained professionals. While threat groups do have some predictable behaviors, individual operators are criminals who may act erratically, causing dire consequences for the victim organization. Engaging qualified professionals gives organizations a clear understanding of threat actor playbooks and current behavior, enabling them to distinguish routine bluffs from credible threats. Expert legal counsel is also essential for understanding reporting requirements and potential legal ramifications. A mature and up-to-date incident response plan can assist with the coordination of all these factors.

Check In: Akira and Qilin Payment Rates

Overview

In our 2025 annual report, we tracked Akira and Qilin as two of the most active and financially impactful ransomware operations of the year.

Akira's 2025 performance was substantially campaign-dependent, driven by mass exploitation of a single vulnerability rather than sustained operational capacity. First half (H1) 2026 has highlighted the group's dependence on such campaigns for operational volume. The group's victims decreased substantially without similar follow-on campaigns.

Conversely, in 2025, Qilin benefited from the diversity of attacks carried out by its affiliates. GRIT observed a wide range of attacks attributed to the group, likely driven by its growing affiliate base. While mass exploitation campaigns lead to greater peaks of activity, Qilin's "slow and steady" approach make them a reliable player in the space.

With H1 2026 now complete, we are revisiting both Akira and Qilin to assess how their payment activity has evolved. The data tells markedly different stories for each. One group is in apparent decline, while the other is holding steadier ground despite a significantly smaller victim pool. Taken together, their trajectories illustrate a core dynamic in ransomware economics: mass initial access campaigns inflate payment volume temporarily, but repetition cannot be guaranteed. Sometimes lightning only strikes once. When discovered and weaponized, broad vulnerability exploitation campaigns can allow experienced ransomware operators to easily infiltrate networks at scale, before capitalizing on that access to commit impactful attacks with payment as the only recovery option for many victims. In the absence of such impactful exploits, payment figures regress along with overall victim volume, and threat actors are unable to profit at the same scale.

Akira

Akira's 2025 payment figures were not a reflection of broad, unconditional operational strength. They were the product of a specific, time-limited access campaign. H1 2026 figures confirm that distinction. Akira's H1 2026 figures reflect a significant YoY contraction compared to H1 2025. Akira's payment volume dropped from 78 payments totaling over $42 million in H1 2025, to just 39 identified payments totaling $8.7 million in H1 2026.

This marks a 50% reduction in payment count and a nearly 80% reduction in total proceeds. The overall reduction is also coupled with a steep decline in the average payment amount, with average H1 2026 Akira payments falling to $223,000 from H1 2025's $413,000. To understand why Akira's economic downturn has occurred, it is important to examine what drove Akira's 2025 performance in the first place.

Akira's elevated 2025 activity was substantially enabled by the mass exploitation of SonicWall devices via CVE-2024-40766, a critical access control vulnerability in the company's SonicOS SSL VPN. The flaw was patched by SonicWall in August 2024 and publicly disclosed in September 2024. Initial exploitation by ransomware affiliates began almost immediately. Akira threat actors harvested credentials from vulnerable SonicWall devices during the initial exploitation window in 2024. It continued weaponizing those stolen credentials well into 2025, even against patched devices, making this campaign particularly effective. During this time, Akira was able to successfully authenticate against MFA-protected VPN accounts using previously stolen one-time password (OTP) seeds. This allowed Akira to enter what were assumed to be protected environments, leading to hundreds of intrusions and subsequent ransomware attacks.

GRIT has not observed a comparable mass exploitation campaign from Akira in 2026. Developing or acquiring the capability to exploit a vulnerability at this scale requires meaningful resources. That could be technical expertise to identify and weaponize a flaw, or capital to purchase an exploit from a third-party broker. The absence of a readily available zero-day to exploit is likely a large contributing factor to Akira's current reduced operational tempo. Akira's hallmark year of 2025 may have simply been an outlier to the group's operational tempo. The group is proving unable to operate at the same scale in 2026 and remains unlikely to do so barring any new, impactful mass exploitation campaigns.

Qilin

Unlike Akira, Qilin did not benefit from an equivalent mass access campaign in 2025. Its payment figures reflect steadier but more modest operational output, which makes the comparison between the two groups instructive rather than incidental.

Qilin's H1 2026 figures show a more moderate shift. Payment count fell from 65 payments in H1 2025 to 43 in H1 2026, which signals a slight volume reduction. However, the average of each payment increased slightly, from $413,000 to $463,000, resulting in a total of $19.9 million in observed payments through H1 2026 against H1 2025's $26.8 million. While on the surface it may appear as though Qilin's success has remained somewhat stagnant based on payment figures, comparing their payment rates to their number of claimed victims reveals a dramatic decrease in effective attacks resulting in payment. Qilin claimed 326 victims to their DLS during H1 2025, which is dwarfed by their H1 2026 count of 659.

Despite Qilin's total payments seeing a slight decline, if they were operating at least the same efficiency, we would have expected their payment figures to increase. Despite conducting presumably more attacks, the attacks themselves may not consistently be impactful enough to warrant higher rates of payment. Through our experiences in working directly with victims of ransomware incidents, we've seen that victims are more willing to pay ransom demands only when absolutely necessary, when backups have been destroyed, and settlement is the only road to recovery. If threat actors cannot destroy backups or achieve complete network encryption, then victims today are less likely to pay, and threat actors are less likely to profit.

One structural factor potentially contributing to Qilin's vast number of affiliated ransomware attacks is their approach to affiliate attraction and retention. Qilin operates as an open-recruitment RaaS, advertising for affiliates on Russian-language cybercriminal forums such as RAMP.

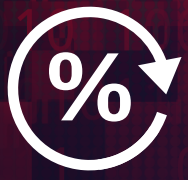

Affiliates receive between 80–85% of ransom proceeds.

This split, while competitive, faces pressure from other groups offering more favorable terms.

Despite this, the group may have an easier time recruiting affiliates, due to what we assess to be a decreased barrier of entry when compared to other groups, including Akira. Qilin requires a monetary deposit to join, whereas Akira's entry likely requires references and personal ties prior to joining their ranks.

For reference, following the disruption of BlackCat/ALPHV in late 2023, that operation announced it was raising affiliate payouts to 90% to retain talent, which signaled how competitive the affiliate market had become. Our assessment in our 2025 annual **Ransomware & Cyber Threat Report** is that affiliate movement between groups following law enforcement action may have pushed displaced affiliates from LockBit and BlackCat towards other RaaS groups, including Qilin. The reverse is also possible: Qilin could lose experienced affiliates to newer or more generously structured operations, resulting in less experienced affiliates remaining who are unable to enact sufficiently impactful attacks.

Implications

Akira and Qilin H1 2026 payment trajectories reinforce a broader market dynamic documented by Coveware: ransom payment rates are declining YoY, and the overall success rate of cyber extortion is contracting. Both groups reflect that trend, but for distinct operational reasons.

Access is the short-term determinant of ransomware payment volume. Groups that achieve mass initial access through vulnerabilities in widely deployed technologies, such as SonicWall, Fortinet appliances, or Check Point VPN, can dramatically increase victim volume, and corresponding payment volume, across the quarters following the campaign. When patch adoption reaches scale or the exploit window closes, payment figures regress accordingly. This makes raw payment data a lagging indicator. H1 figures typically reflect campaigns executed months prior, not current operational capacity.

Akira's decline is the clearest illustration of this. A group that posted $42 million in H1 2025 on the back of a single CVE has contracted sharply in its absence. Qilin's ability to conduct a wider scale of attacks, by contrast, reflects a structural model less dependent on any single access vector, which may prove more durable over time, even if it produces less payments compared to victim volume. The question for H2 2026 is whether Akira reestablishes access at scale, and whether Qilin's affiliate base holds.

Note: *We want to thank our friends at TRM labs for access to the information to conduct this analysis.*

Industry Spotlight: Banking and Finance

The defining risk of ransomware in the banking and finance sector is the industry's dependence on customer confidence. Banks, investment firms, rating agencies, and market infrastructure providers must preserve service availability, data confidentiality, and trust in the integrity of their underlying data. Customers in this industry need to feel extremely confident that their private information, access to money, and that their financial future are secure. Ransomware actors that target this industry seek to encroach upon this confidence to extract extortion payments. The **FBI's Internet Crime Complaint Center (IC3) Internet Crime Report** recorded a combined total of 7,574 ransomware and data breach complaints in 2025. 447 of those involved financial services organizations and making financial services the third most attacked sector.

GRIT's quarterly tracking of ransomware victims shows that the pace of victims amongst financial services firms appears to have modestly declined 16.18% YoY, with 57 observed incidents in Q2 2026 compared to 68 in Q2 2025. Our data also shows that the financial services sector is no longer among the top ten most impacted industries in Q2 2026. Despite this, the sector's exposure extends well beyond direct attacks.

Financial institutions sit at the center of interconnected service ecosystems. A breach anywhere in the supply chain frequently leads to impact on financial institutions. As our understanding and approach to supply chain attacks and downstream impacts continues to develop, incidents historically viewed as purely technical setbacks can now be considered as potentially destabilizing institutional trust. For the financial services sector in particular, resilience to ransomware threats is transitioning into a defining measure of operational credibility.

Due to government regulation and oversight, the financial services sector often has more significant security, governance, and resilience requirements than in other industries. At the same time, increased regulatory oversight also means ransomware incidents impacting the financial industry carry legal, supervisory, and reputational consequences that extend far beyond the initial technical response.

- Under 12 CFR Part 53, banking organizations must notify their primary federal regulator within 36 hours of determining a notification incident has occurred (OCC/FDIC/Federal Reserve, Final Rule, November 2021).
- Concurrently, publicly traded financial institutions must disclose material cybersecurity incidents via SEC Form 8-K (Item 1.05) within four business days of a materiality determination (SEC Release No. 33-11216, July 2023).

Together, these overlapping requirements ensure that a single ransomware event triggers parallel regulatory, legal, and public disclosure timelines. This compounds pressure on incident response teams already managing operational recovery.

Similarly, the **OCC's 2026 Cybersecurity and Financial System Resilience Report** treats financial institutions and key third-party providers as part of the same regulatory problem. Strong regulations like these improve the baseline preparedness of an institution but also multiplies their obligations in the event of a breach. Different teams within the organization being attacked must simultaneously determine scope, maintain operations, satisfy reporting obligations, and manage public confidence. The increased burden of a fully compliant response can make improvised response in itself a material risk to both the institution and its customers.

Regulatory controls have not fully eliminated the risk posed by ransomware and extortion, especially across private markets and digital finance. In response, law enforcement has expanded sanctions, seized illicit infrastructure, and targeted cryptocurrency exchanges used to facilitate ransom payments and laundering.

In March 2025, the DOJ disrupted Garantex, freezing more than $26 million in allegedly illicit assets, while the Treasury Department linked over $100 million in Garantex transactions to ransomware groups and other sanctioned actors. Despite these efforts, private equity and credit, fintech, and digital asset firms remain exposed to supply chain attacks through increasingly interconnected technology ecosystems, where shared service providers, cloud platforms, APIs, and identity systems create systemic potential points of failure. An incident affecting one administrator, portfolio company, or technology provider can cascade across multiple funds, counterparties, and investors, extending well beyond a singular compromised organization.

The financial services sector's cybersecurity posture is facing a conceptual gap. Most organizations still define resilience as system restoration, but data theft is quickly becoming a significant targeted threat. Once confidential market-sensitive data is exfiltrated, the attacker retains leverage independent of whether systems are fully recovered. That exposure extends beyond direct attacks. Financial data can be compromised when threat actors hit the downstream law firms, benefits administrators, and service providers that hold it on behalf of clients. Resilience in this sector must mean protecting information that, once exposed, cannot be 'recovered' at all.

Threat Actor Spotlight: FulcrumSec

FulcrumSec is a data extortion group first observed “in the wild” in October 2025. That said, they did not begin regularly claiming victims until late April 2026. FulcrumSec is distinguished by its use of pure data extortion, focusing on exfiltration of large quantities of victim data, particularly from cloud environments, while eschewing encryption. The group has adopted a heavily westernized lexicon in its communications and data leak site, referencing “The Hardcoded Horror Show”, the “Index of /Shame,” and “SLOPOCALYPSE NOW” as breach victim types. They sometimes employ British English spelling in their communications, suggesting the possible participation of English speakers as part of the organization. This mirrors a trend we have seen of cloud-based data theft carried out by ShinyHunters, LAPSUS$, and other groups tied to “The COM” in recent years; however, we cannot definitively tie FulcumSec to this origin.

FulcrumSec has repeatedly detailed how it achieved access to victims during negotiations. GRIT has analyzed to gain insight into the group's targeting, tactics, techniques, and procedures. In reviewing the group's claims across over 13 victims, we observed the following repeated details:

- FulcrumSec operates seemingly entirely within cloud provider APIs and management planes. The group enumerates and exfiltrates data at scale from AWS S3 Buckets, Google Cloud and Azure platforms, as well as Firebase and MongoDB.
- FulcrumSec gains initial access by exploiting exposed, misconfigured, or stolen cloud credentials rather than deploying traditional malware or phishing lures against end users. Across all identified victims, the group accessed cloud environments through service account keys, IAM credentials, OAuth2 RSA private keys, Firebase tokens, and API keys. In several cases, they sourced these credentials directly from application logs and configuration files left readable within the victim environment. FulcrumSec appears to conduct systematic post-access enumeration of source code, configuration files, and application logs specifically to harvest live credential material.
- FulcrumSec has consistently collected personally identifiable information (PII) and protected health information (PHI) and intellectual property (IP) as part of its data exfiltration focus. The group also appears to favor broad downloads of databases, which they have at times joined to form a more useful and holistic compromise of personal information in aggregate.
- In addition to their Data Leak Site, FulcrumSec has posted victim data and breach details to the illicit forums DarkForums and BreachForums since October 2025. The group often emphasizes the sensitivity of victim's allegedly compromised data while soliciting victims, researchers, and journalists to contact them for additional details.
- The group appears to use AI/LLMs heavily to parse large volumes of data and summarize key findings of interest or sensitivity, both in victim communications and on its data leak site. The group's authors repeatedly describe breaches as "the worst they've ever seen" or "a disgrace". They use similar emotional language while describing victims. On at least one occasion, the group refused to publish allegedly breached data due to the impacts the leak would have on "innocent individuals", framing themselves as magnanimous.

In summary, FulcrumSec represents another in a growing line of data-extortion-focused, seemingly westernized threat actors in the mold of ShinyHunters, LAPSUS$, and Scattered Spider. They continue to prove that encryptor deployment is not a prerequisite for receiving a ransom. Additionally, without a core franchise building and maintaining an encryptor, there are less kickbacks required out of any ransoms received.

The uncomfortable implication of this developing approach is that the growing recovery capability enjoyed by many well-defended organizations that consists of redundant, immutable, segmented backups, which does not and cannot prevent wholesale data extortion, and the ability of an organization to recover from an exfiltration-only attack does not influence the demand of data extortionists.

Defenders can best prepare for this category of threat through:

- Continuous secrets scanning
- Automated rotation enforcement for API keys and service account credentials
- Eliminating long-lived credentials in general wherever possible
- Implementing dark web monitoring or similar tools and services to detect credential compromise before it can be weaponized

Phishing-resistant MFA and help desk authentication protocols remain paramount for the social engineering aspects of these attacks, and an increased focus on Cloud security becomes a necessity. Organizations with any cloud-based network components should seek to detect publicly exposed storage buckets and databases. They also should enable robust logging and alerting which are often disabled by default and apply conditional access policies and network segmentation to restrict unauthorized access with valid credentials and further lateral compromise.

Quarterly Wrap Up

Other than the normal ebbs and flows, trends in the ransomware ecosystem in Q2 2026 have played out as anticipated. Despite an influx of new players, most victims claimed this quarter were from known groups responsible for a disproportionate share of victims. While numbers are still slowly rising, we did not see as significant of a jump as we have had in more dramatic quarters (e.g. Q4 2025). Similarly, forecasts of widespread novel AI deployment by cybercriminals have not yet come to fruition, with AI adoption similarly following a slow but steady and predictable increase over time.

It seems perverse to call something as malicious as the ransomware ecosystem “healthy,” but the landscape is now marked by multiple entrenched players who show no sign of slowing down. Going into Q3, we typically expect a small decrease in observed victim volume, perhaps driven by summer holidays or rising temperatures. It will be interesting to observe if the current, top-heavy, and efficient structure of ransomware operators can push through this recurrent seasonal decline.

Yet this persistence comes with a cost. Threat actors are being forced to adopt to increasing Defender preparedness. While 2025 reflected heavy attacks on edge devices, hardened perimeters now face threats to identity, cloud, and data exfiltration, which can evade newly placed safeguards. These approaches often entail greater scale, which offsets their associated costs and lead-time. Ransomware figures are up, and these increases are occurring despite heightened defenses, and as threat actors adopt new techniques.

Opportunistic actors are constantly looking for ways to compromise new victims while simultaneously getting better at turning small compromises into big ones. Given the success of recent attacks, we can expect to see at least one more large-scale supply chain attack before 2026 ends. As defenders harden their networks, attackers continue to refine their techniques. Thus, the cat and mouse game continues.

We encourage our fellow Defenders to recognize and appreciate the improvements they have made in response to last year’s threats, and to recognize that this will provide greater security – until it doesn’t.

As always, it is imperative that we observe these latest tactics with an eye towards defense, mitigation, and detection, even as they remain in flux.

Happy Hunting!

Methodology

Data collected for this report was obtained from publicly available resources, including threat groups themselves. It has not been validated by alleged victims. Collected data is reviewed for potential duplications or inaccuracies and are adjusted accordingly. Thus, the number of publicly observed attacks and the actual number of attacks conducted may not be equal. Some groups do not publicize all their victims, and almost all groups offer an option to withhold announcement if the victim pays a ransom within a specified timeframe and/or remove the victims once a ransom has been paid. Additionally, some groups include incomplete information about their victim or claim an attack despite successfully attacking only a small subset of their target. For these reasons, the data in this report is useful in aggregate, but should be evaluated as a report consisting of data sources that have variability. Despite the variability, this report is still an accurate representation of the total ransomware threat landscape.

We note that this report includes data and analysis of several groups that may be better described as “extortion” groups rather than “ransomware” groups. These groups may eschew encryption and instead focus only on data exfiltration and extortion, or may not perform intrusion operations of any kind, instead extorting or re-extorting organizations based on historically compromised data. While these groups do not deploy ransomware, we are including them in our reporting due to their relationships with other ransomware groups and their impact on the extortion-based cybercrime environment.

Finally, we make efforts to exclude from our data those groups that self-identify as “hacktivists”, compromised data brokers and markets, or non-financially motivated data thieves and leakers. While these actors and venues doubtlessly have impacts, we distinguish them from financially motivated cybercrime and data extortion which is the primary focus of this report. For this reason, our data may periodically reflect lower total numbers of incidents than other, similar public reports.

1900 Reston Metro Plaza, Suite 701, Reston, VA 20190
guidepointsecurity.com • info@guidepointsecurity.com • (877) 889-0132

R.GRITQ2.2607