

A GRIT® REPORT

Ransomware and Cyber Threat Insights

Q3 2026

July–September



GUIDEPOINT®
SECURITY

Table of Contents

	Quarterly Ransomware Summary	3
	Threat Actor Spotlight: ShinyHunters	4
	Industry Spotlight: Education	5
	Vulnpocalypse: The Sky Isn't Falling, But the CVEs are Piling Up	6
	Threat Actor Spotlight: Clop is Flopping	9
	Ransomware Payments Are Down. So Why Isn't Anyone Relaxing?	12
	Tinker, Model, Agent, Spy: How Ransomware and Cybercrime Groups Are Using AI in 2026	14
	Quarterly Wrap Up	18
	Methodology	19




Quarterly Ransomware Summary

Welcome to the Q3 GRIT Ransomware and Cyber Threat Insights Report. This quarter, we look at the so-called “vulnpocalypse,” shine a spotlight on the Education industry and ShinyHunters group, analyze the Clop syndicate’s latest activity, and assess payment trends in the ransomware space. As always, we begin with a summary of the quarter’s ransomware activity, which shows no sign of slowing.

In Q3 2026, the GuidePoint Research and Intelligence Team (GRIT) tracked 2,760 victims claimed by 112 distinct threat actors across 29 industries and 115 countries. These figures represent a significant Quarter-over-Quarter (QoQ) escalation over Q2 2026. Q3 saw a 21.1% QoQ increase in total victims and a 23.1% rise in the number of active threat actors, while the geographic spread of targeting expanded by 6.5% from 108 to 115 countries. The number of targeted industries remained effectively flat, declining marginally (–3.3%) from 30 to 29.

Year-over-year (YoY) comparisons underscore the accelerating pace of ransomware operations. Relative to Q3 2025, victim counts surged 75.3% YoY (up from 1,574), the number of active threat actors grew by 47.4% (up from 76), and targeting expanded to 25 additional countries (a 27.8% increase), while industry breadth held steady. Taken together, these trends indicate that the ransomware ecosystem continues to grow in both scale and operational diversity, with a broadening threat actor landscape driving record-setting victim volumes.

Q3 2026 Rate of Publicly Posted Ransomware Victims per Week



	Q3 2026	Q3 2025	Q2 2026
Total Publicly Posted Victims	2,760	1,574	2,279
Active Ransomware Groups	112	76	91
Average Posts Per Week	197	113	175
Average Groups Posting Per Week	41	25	33

On the ransomware group front, TheGentlemen usurped Qilin as the most active group during Q3 2026, but only by a small fraction, as both groups dominated, each with 13% of the market share. Traditional heavyweights like Akira and Dragonforce took a back seat to Inc and Krybit, with Clop, Safepay, Storm, and Direwolf rounding out the top 10. Despite sanctions, Lockbit continued to appear in the top 20 active groups, with relative newcomers crpxo, orova, and panzer also making the leaderboard. While the newcomers are making their way up the charts in Q3, the trends we observed in the threat landscape maintained the momentum we saw in Q2, with only a minor QoQ shifting of the major players among the top ten.

What does this mean for the ransomware ecosystem and for defenders? Largely, we still have not hit peak ransomware, and the threat is real, though perhaps not directly as financially profitable as it has been historically. However, the burden of effort for defenders has increased, and the technical and administrative fallout from the effects of a ransomware attack can be just as detrimental to business functionality as making a payment would be. As 2026 moves into the last quarter of the year and 2027 looms, we will continue to monitor the ransomware ecosystem.



Threat Actor Spotlight: ShinyHunters

ShinyHunters is now one of the most active and visible extortion groups of 2026. The group illustrates how professionalized cybercrime groups can become. First observed in 2019, the group has matured into an operation that hits large, high-value targets, adapts quickly when defenders respond, and keeps itself in the headlines. In Q2 and Q3 this year, it ran a mass exploitation campaign that compromised more than 100 organizations, claimed a breach of the FBI, and hacked rival ransomware gang Clop.

ShinyHunters is a “pure” data extortion actor, meaning it skips ransomware encryption entirely, steals data in bulk, and demands payment to keep it private. This operating model lends itself to scale. With nothing to decrypt and no recovery to negotiate, each operation is simpler to run and easier to repeat. This mirrors a broader trend GRIT has tracked across the cybercriminal ecosystem, where actors are moving away from double extortion. The model also pays. Tracked payments attributable to ShinyHunters average approximately \$600,000 per victim in 2026, with more than \$8 million USD observed across attributed Bitcoin wallets since the beginning of Q3 2026.

The June 2026 exploitation of CVE-2026-35273, a zero-day vulnerability in Oracle PeopleSoft’s Environment Management Hub, shows that the pure extortion model works in practice. The initial wave alone compromised more than 100 organizations across education, healthcare, government, and technology. When Oracle patched and defenders deployed web application firewall rules to block the vulnerable endpoint, ShinyHunters encoded a single character in the request path (/ %50SEMHUB/ in place of /PSEMHUB/) to bypass those controls. Within weeks, actors resumed mass exploitation. From there, the group deployed web shells, a custom backdoor (SIDEYE), and tunneling tools to move deeper into victim networks.

In late September, ShinyHunters claimed to have breached the FBI, which would be its most prominent target to-date. The group claims to have accessed

FBIjobs.gov through the same PeopleSoft vulnerability. They claim to have exfiltrated 2–3 terabytes of data, including personnel records, psychiatric evaluations, fitness-for-duty examinations, and medical files on active agents. [Reuters](#) has partially authenticated several of the documents, and the FBI has publicly confirmed an investigation. Regardless of whether the full scope of the claim is substantiated, the data allegedly involved (e.g., medical health records, counter intelligence assignments, and medical details) is likely to attract interest beyond financially motivated actors, including foreign intelligence services.

Between the PeopleSoft campaign and the FBI claim, ShinyHunters also compromised the Clop ransomware gang’s data leak site in mid-September 2026, defacing it, stealing server logs and Onion service (TOR stack) private keys, and issuing an extortion demand against Clop itself. The group says the attack was retaliation for threats a Clop representative made during a dispute over the Oracle E-Business Suite campaign. The episode matters less for the drama than for what it shows: ShinyHunters is confident enough to attack a peer, and it appears to welcome the attention that follows.

The group’s high-profile activities may be catching up to them, however. [Dutch police](#) confirmed the September 15, 2026, arrest of an alleged ShinyHunters leader. [Reuters](#) reports that a second alleged member was detained in Jordan on September 29 and is cooperating with the FBI. Regardless, ShinyHunters enters the final quarter of 2026 as arguably the most active pure extortion actor in the threat landscape, with the reach to credibly threaten U.S. government infrastructure, the skill to adapt to defenses in near-real time, and it appears uninterested in maintaining a low profile. The questions now are whether law enforcement pressure broadens into coordinated action that disrupts the group’s operations, whether the allegedly stolen FBI personnel data surfaces in secondary markets or foreign intelligence channels, and/or whether the group’s appetite for escalation draws retaliation from Clop or affiliated actors.



Industry Spotlight: Education

The education sector remains a frequent target for ransomware attacks, with both schools and education service providers at risk of experiencing significant operational, financial, and reputational impacts. This industry spotlight examines the evolving ransomware threat to education, highlighting key trends, attack patterns, and the challenges organizations face in protecting critical systems and sensitive student and institutional data. Educational facilities, as much as any other industry, no longer run on paper and most rely on software as a service (SaaS) offerings for priority tasks from scheduling to assignments and materials. By their very nature, these platforms store personally identifiable information relating to minors, which unfortunately makes them a target for financially motivated threat actors seeking to steal data for the purpose of extortion.

This dynamic increasingly extends beyond individual schools to the SaaS vendors that support them. In May 2026, ShinyHunters claimed a ransomware and extortion attack against Instructure, the U.S.-based parent company of the widely used Canvas learning management system. The attack disrupted access to coursework and grading systems for many institutions during a period of the spring term when reliance on the platform was at its highest. Over the following months, GRIT tracked a further six education-technology and learning-management platforms targeted in the same manner, spanning multiple countries and claimed by six distinct threat actors working independently of one another.

GRIT assesses this pattern reflects convergence on a shared target logic rather than a coordinated campaign. Compromising a single vendor can expose the data, and in some cases the access, associated with every institution it serves, all for the opportunity cost of one intrusion. For attackers, that math is generally more favorable than pursuing schools individually. For the institutions themselves, it means meaningful risk exposure can originate well outside their own network perimeter, in systems they neither own nor directly control.

Separate from the vendor layer, the accounts belonging to students and staff are themselves a recurring point of entry. Multi-factor authentication adoption remains inconsistent across institutions, password hygiene is often weak, and broad system access is routinely granted to a population that receives little to no security training. Once compromised, GRIT assesses that these accounts can serve as a foothold for further social engineering or lateral movement toward more sensitive institutional systems or organizations.



The average K-12 school in the United States allocated **8%** or less of its IT budget to cybersecurity, while **18%** of schools allocated less than **1%**.

This same environment also warrants consideration from an insider risk perspective. Legitimate account holders already possess valid access, so misuse does not necessarily require a full compromise, and the permissive conditions that make these accounts easy to steal make them equally easy to misuse directly.

The underlying conditions in both cases are structural rather than incidental. Enrollment turns over completely within a few years, accounts are provisioned quickly with minimal identity verification, and a “.edu” credential often carries more implicit trust than the vetting behind it warrants. Accounts belonging to students who have graduated, transferred, or withdrawn frequently remain active well after they stop being monitored, extending the window of exposure long past a student’s period of active enrollment.

Underlying both dynamics is a persistent funding gap with implications that extend beyond technical posture. Most K-12 districts lack the budget to build a mature cybersecurity program. According to a [2022 study by the Center for Internet Security](#), the average K-12 school in the U.S. allocated 8% or less of its IT budget to cybersecurity, while 18% of schools allocated less than 1%. Many chronically underfunded school systems still spend effectively nothing on cybersecurity and rely on built-in tools and security features which crumble when interrogated by advanced threat actors.

Ultimately, the Education sector faces a challenging combination of an expanding SaaS attack surface and limited funding for cybersecurity protections. As schools increasingly rely on cloud-based platforms and third-party services, gaps in security controls and resources can create opportunities for ransomware and other financially motivated threat actors to gain access, steal sensitive data, and potentially disrupt critical operations. This risk is unlikely to go away any time soon, and digital integrations in schools may continue to present more risk. Now more than ever, schools must focus on identity and access management security in addition to the traditional network perimeter. Even though budgets are tight, investment in cybersecurity should be prioritized by these institutions who are responsible for safeguarding our most vulnerable population’s data.




Vulnpocalypse: The Sky Isn’t Falling, But the CVEs are Piling Up

There has been a dramatic uptick in the number of Common Vulnerabilities and Exposures (CVEs) published across major organizations over the past several months. With 978 CVEs recorded in Microsoft’s Patch Tuesday in September alone, the trend warrants closer examination. This trend holds true across major vendors like GitHub, Apache, Mozilla, and Microsoft, but for the purposes of this section, we’re narrowing our lens to [Microsoft’s Security Updates](#).


What’s driving this surge? Is this consistent with historical patterns? Are these vulnerabilities seeing active exploitation in the wild? What contributing factors explain the spike? And, perhaps most importantly, what does this mean for the defenders tasked with triaging and remediating them?

Microsoft Security Updates CVEs



Compared to previous years, the surge in vulnerabilities represents a marked departure from established trends. The graph above tracks Microsoft Security Update CVE volume over the past three years. Across 2024, 2025, and 2026, vulnerability disclosures followed a comparable trajectory in the first half of each year. That changed in July 2026, when 663 disclosures were recorded compared to just 144 and 137 in the same month in 2024 and 2025, respectively. That gap has continued to widen in subsequent months. The most probable driver behind this acceleration is the growing integration of AI into vulnerability discovery and security research workflows, which will be examined in detail below. First, however, we will assess the impact of these disclosures.

Microsoft CVE Exploitation Potential in 2026



While the volume of disclosed CVEs has surged dramatically, the distribution of severity ratings has remained largely unchanged. As part of each Security Update, Microsoft categorizes vulnerabilities by exploitation likelihood: detected, more likely, less likely, or unlikely. As the graph above illustrates, most CVEs published throughout 2026 have been classified as “less likely.” Aggregating the less likely and unlikely categories, 86% of disclosed vulnerabilities fell into these lower-risk classifications in 2024, 83% in 2025, and 83% through 2026 to date. This shows that while the raw volume of disclosures has surged dramatically, the proportion of high-risk, exploitable vulnerabilities has remained flat, suggesting the overall quality of findings has declined alongside the increase in volume.

The most likely catalyst behind the surge in disclosed vulnerabilities is the increasing adoption of AI across security research and tooling. It’s worth noting, however, that many of these disclosures have not been explicitly credited to AI-assisted discovery. It remains unclear whether that’s because organizations don’t feel compelled to disclose their use of AI in the discovery process, or because AI has become so deeply embedded in modern security tooling that it no longer warrants a separate callout. Regardless, the practical implication for defenders is clear: the upward trend in CVE volume is likely here to stay. The rising tide of disclosures doesn’t, however, necessarily translate to a proportional increase in weaponizable vulnerabilities, at least not within the scope of Microsoft’s published Security Updates. The signal-to-noise ratio remains manageable for now. The challenge will be in ensuring that triage processes and patch management programs scale accordingly.



Threat Actor Spotlight: ClOp is Flopping

ClOp (sometimes stylized as “ClOp”) is a long-running ransomware and data-extortion operation distinguished by its exploitation of zero-day and newly disclosed vulnerabilities in widely used enterprise file-transfer and back-office software. The group uses the vulnerabilities to compromise victims at scale and extort them through mass data theft rather than on-network encryption. The group found relative success in 2023 with large scale campaigns exploiting Forta’s GoAnywhere MFT and Progress Software’s MOVEit file transfer solutions. The campaigns netted the group tens of millions of dollars in ransom payments, as per [our previous analysis](#) on the group’s activities.

ClOp has been unable to recreate previously seen success with subsequent mass exploitation campaigns. Their newest endeavor targeting PTC Windchill has not resulted in a terrible amount of success. We have not identified any ransom payments made to ClOp from PTC Windchill campaign, and the volume of victims shared to their data leak site (71 at the time of writing) is dwarfed by the number of victims from previous campaigns, including over 250 and 350 from MOVEit MFT and Cleo MFT exploitation campaigns, respectively.

ClOp’s unsuccessful campaign was also followed by a public dispute with the data extortion group ShinyHunters. ShinyHunters claims that their attack, resulting in the defacement of ClOp’s dark web site, is the result of threats made by ClOp that were directed at ShinyHunters. ShinyHunters is now threatening to leak information pertaining to victims who paid ClOp (and ostensibly paid for ClOp’s silence due to data theft), as well as the respective Bitcoin wallet addresses used for the payments.

The compounding effect is the real problem. A low-yield campaign hurts ClOp financially. But the ShinyHunters dispute threatens ClOp’s credibility, which is arguably more valuable than any single campaign’s proceeds. Ransomware and data-extortion business models depend entirely on victims believing payment buys silence. A public spectacle

over which victims paid and how much they paid directly undermines that trust, both for this campaign and every future activity. Even a technically successful campaign going forward may yield fewer payments if victims no longer believe ClOp will keep its word. ClOp is already a shell of their former selves, akin to a star racehorse in their elder years that may need to be taken out to pasture in retirement. Absent a rebrand or major shift in tactics, ClOp’s operations under their current brand may be close to end of life.

PTC Windchill Campaign

ClOp’s most recent campaign targeted PTC Windchill and its retail add-on FlexPLM, a product lifecycle management (PLM) platform. The exploited vulnerability, CVE-2026-12569, is a critical, unauthenticated, improper-input-validation flaw that enables remote code execution via threat actor requests. [PTC disclosed the vulnerability](#) on June 17, 2026, but reporting from Censys indicates exploitation began as a true zero-day in early June, preceding the public advisory. PTC has since disclosed two related vulnerabilities (CVE-2026-77645 RCE and CVE-2026-77646), and it continues to publish new IOCs and webshell-path patterns through at least late August 2026, indicating the intrusion set is still active.

ClOp began sharing redacted victim names accrued from the PTC Windchill exploitation campaign on its leak site on August 12, 2026. They pivoted to publishing full company names, and by mid-to-late August, ClOp had named more than 40 organizations. Numbers grew to roughly 71 distinct victim postings by late September 2026. Windchill has a notably smaller attack surface than prior ClOp mass exploitation targets. Censys observed fewer than 100 exposed Windchill/FlexPLM instances globally as of June 1, 2026, a figure corroborated by a live Shodan query returning 110 internet-facing results. By comparison, MOVEit Transfer had over 2,000 exposed instances and Oracle EBS over 2,700 during their respective ClOp campaigns. ClOp’s scope of potential

targets is notably smaller with the newest campaign, which is likely an attributing factor to the decrease in both targets and successful extortion payments.

Several companies named by Clop have publicly acknowledged investigating the claims, but none have confirmed material breaches at the time of this report. We assess with moderate confidence that no ransom payments have been publicly confirmed for this Windchill/FlexPLM campaign to date. Much of the stolen PLM data (engineering files, product specs, facility images) appears to be intellectual property rather than regulated customer Personally Identifiable Information (PII), giving victims comparatively less regulatory pressure to pay.

Data Extortion Can Be Impactful If the Data Matters

Data extortion without encryption can be extremely lucrative, but typically only when the access vector lets the attacker choose targets with guaranteed high-value data. ShinyHunters and Silent Ransom Group (SRG) demonstrate this as two prominent data-extortion groups that forgo encryption in their attacks. The groups rely on targeted social engineering attacks rather than technical exploitation of software vulnerabilities. SRG has pivoted from law firms to hedge funds likely because those targets are known in advance to hold dense concentrations of sensitive financial and legal data. ShinyHunters' Salesforce/Salesloft Drift campaign similarly went straight at CRM platforms, guaranteeing access to customer PII and business records. Because the access method is human-directed rather than exploit-driven, these groups can be selective, low-noise, and consistently effective at extracting negotiable leverage. ShinyHunters and SRG's disclosed victims and reported payment outcomes bear this out when compared to Clop.

Clop's model is the inverse. It is more technically sophisticated with zero-day exploitation of enterprise software. This method allows them to compromise dozens of organizations in a single wave without any social engineering or target research. But this scale comes at the cost of precision. Clop has no control over what data resides on a given victim's instance of the exploited software. In the Windchill campaign, that data presumably is largely PLM/engineering material. Although it may be valuable as intellectual property,

it is not the kind of urgent, regulatory charged, or immediately monetizable data that drives fast ransom payments. The result is visible in the outcomes. Despite more than 80 named victims, we have not confirmed any payments from this specific campaign, with even high-profile targets publicly stating they are "investigating" rather than negotiating.

Mass technical exploitation produces volume, but not an equal return in value for Clop. Other data extortion-only groups have demonstrated that social engineering can produce value, without the need for additional technical debt. Clop's approach is a numbers game wherein they cast a wide net via a single software flaw and hope some fraction of victims have data worth paying to protect. This is likely why Clop has increasingly sought collaboration with social-engineering-driven collectives rather than reverting to on-network encryption. Flashpoint has reported that in Clop's 2025 Oracle E-Business Suite campaigns, Clop's exploit access was paired with (or overlapped with) "Scattered LAPSUS\$ Hunters," a coalition drawing on ShinyHunters/Scattered Spider-style vishing and identity-abuse tradecraft. This collaboration suggests Clop recognizes that its own technical access needs to be paired with better target/data selection to convert into payment. Rather than a return to double extortion, the more probable evolution is Clop increasingly outsourcing or partnering on the social-engineering side of the operation while retaining its core competency mass exploitation of enterprise software as the initial access method. This would be the case, however Clop's collaborators may not be willing to continue any partnerships.

ShinyHunters Turns on Clop: Extortion Within the Extortion Ecosystem

ShinyHunters and Clop's collaborator loyalties are now being tested in public, as ShinyHunters is now ostensibly extorting Clop. ShinyHunters shared updates to their dark web site, in which the group claims to have worked alongside Clop during the 2025 Oracle E-Business Suite (EBS) campaign. Whether the original collaboration was as deep as ShinyHunters now suggests is unverified, but ShinyHunters claims to have knowledge of victims that were attacked during the EBS campaign.

On its data leak site, ShinyHunters posted a direct threat naming individuals allegedly tied to Clop's operation:

[19 Sep]: "IF YOU WANT TO SAVE YOUR BRAND AND NOT DIE BY MY HANDS: Email us from your official email at shinygroup[@]onionmail[.]com and lets see how wealthy you really are. 2.333% of my networth is a 8 figure amount, I hope you can pay that much because that is the demand, negotiable. Get your bosses in front of the white board in the war room. Clock is ticking moron. Kindly excuse our unprofessionalism."

[20 Sep]: "Dear Likhogray & Tarasov, tell your boss j0nny to wake the <EXPLICITIVE> up. Run those pockets. I want all the money you made off the EBS campaign plus more AND WITH INTEREST. before I start releasing information regarding the companies that paid you, how much, and to what Bitcoin address. My phone book contains all major financial media outlets. CLOCK IS TICKING! LETS GET THE BALL ROLLING! Be sure to bring an English interlocutor so you can comprehend my literacy in acquiring your bank account. 66 hours remaining."

[21 Sep]: "Every 24 houts you fail to engage with us the demands increase. The demand now includes a mandatory apology issued directly to me PUBLICLY. I am 3-0 against you rich and broke criminals."

ShinyHunters is claiming to hold specific financial intelligence on Clop's EBS campaign and threatening to disclose that data to "major financial media outlets" if Clop does not pay a ransom demand. ShinyHunters scoped the demand at 2.333% of ShinyHunters' claimed net worth, described as "8 figures."

Clop's public response, posted to its own leak site, was notably muted and procedural rather than combative:

"Shiny Hunters we trying to reach you Your email does not work / Come online old platform no email"

ShinyHunters' approach of naming individuals, setting a countdown, threatening media disclosure, and demanding a negotiable dollar figure is the same pressure playbook it uses against corporate victims, now redirected at a criminal peer. Clop's reply, by contrast, reads like an operational-support ticket, in that it does not deny the claims or counter-threaten. It simply notes a broken email address and asks ShinyHunters to use "the old platform." That asymmetry of aggressive, theatrical, financially-literate pressure on one side, and flat, administrative, almost passive engagement on the other mirrors the broader extortion behaviors between the two groups' negotiation styles.

For threat intelligence purposes, this dispute is worth monitoring closely for several reasons. If ShinyHunters follows through and discloses victim/payment/Bitcoin data from the EBS campaign, it would provide the clearest public confirmation to date of Clop's actual ransom proceeds and paying victims, useful for validating or correcting payment-rate assumptions elsewhere in Clop reporting. Additionally, Clop paying another ransomware/extortion group

to avoid its own financial data being leaked would be a significant reputational and operational blow, potentially accelerating fragmentation, rebranding, or retaliation within the Clop/FIN11 ecosystem. Given Clop's demonstrated preference for low-friction, low-engagement negotiation (consistent with its instructions to victims and its terse reply here), we assess it is more likely to attempt to relocate or go quiet than to publicly capitulate to ShinyHunters' demand. It is important to note that this is an active situation, and any resolution should be tracked as a direct indicator of Clop's financial exposure from the current campaign.

Conclusion

Clop's recent campaigns have proven less impactful both in terms of victim volume and payments received, and the Windchill PTC campaign is no different. Furthermore, Clop has accidentally entered a public dispute with ShinyHunters, which could end poorly for Clop. Clop may not survive in the current ransomware ecosystem and may need to seek an evolution in tactics or rebranding.



Ransomware Payments Are Down. So Why Isn't Anyone Relaxing?

Right now, a popular narrative has taken hold in the cybersecurity industry: fewer ransomware victims are paying, so the tide must be turning. [Chainalysis](#) reported that total on-chain ransomware payments fell 8% YoY to \$820 million in 2025. The [Verizon DBIR](#) noted that 69% of victims refused to pay, up from roughly half just a few years ago. [Sophos](#) found that median ransom demands and payments both dropped YoY. On the surface, this looks like progress; that organizations are getting better at saying no.

However, by nearly every measure that matters, ransomware attacks are accelerating. From Q1 to Q3 of 2026, GRIT tracked 7174 victims among 155 threat actors across 30 targeted industries in 137 countries. Compare this to 5228 victims, 110 threat actors, 30 targeted industries, 121 countries over the same period in 2025, and ransomware has instead accelerated. If payments are down but incidents are up across every measurable dimension, the simpler explanation is not that ransomware is becoming less profitable. It is that it is becoming more widespread, and that fewer victims are choosing to pay.

Fewer Are Paying, but Those Who Pay Are Paying More

The headline statistic that payments are down obscures a different reality. While the rate of payment has declined, the amount paid when victims do pay tells a different story. Chainalysis itself noted that the median ransom payment rose 368% YoY, from \$12,738 in 2024 to \$59,556 in 2025. [Zscaler](#) reported the average ransom payment climbed 5.3% to \$431,995 in its most recent reporting period (April 2025 to March 2026).

Our own data from threat actor negotiation engagements reflects a version of this tension. In Q3 2026, we handled 29 ransomware engagements, a 61% YoY increase from 18 in Q3 2025. The rate of payment dropped sharply, from 50% to just under 21%. But among those who did pay, the average

payment rose 34%, from \$240,000 to \$321,000. The story in Q2 is more complicated. Engagements more than doubled YoY from 11 to 24, while the average payment fell from \$1.2 million to roughly \$498,000. The 2025 and 2026 figures were heavily skewed by seven-digit ransoms across a much smaller sample set. Remove those outliers (including two 7-digit ransoms in 2025 and one 7-digit ransom in 2026), and the underlying per-payment figure trends much differently, from \$58,000 to \$138,000, a 138% YoY increase.

The takeaway is that the relationship between payment frequency and payment size is more complex than the headlines suggest. Fewer victims are paying, but the ones who do are often facing demands and making payments that remain substantial.

The Big Game Hunt Is Evolving

For years, the dominant ransomware strategy was what the industry calls “big game hunting” targeting large, well-resourced organizations capable of paying seven- and eight-figure ransoms. Groups like Conti, LockBit, and BlackCat built their reputations on this model. The question now is whether that playbook still works.

The answer is nuanced. Researchers at [Cardiff University](#) recently characterized the current era as the “Ransom Bazaar”: a “financialized, service-oriented extortion” ecosystem that evolved out of the big game hunting phase. The single massive payday still happens, but it’s no longer the primary engine of the economy. Instead, the landscape has fragmented. The barrier to entry has dropped, and the affiliate model means that a wider pool of operators are running campaigns across a broader target set.

Our data hints at this shift. As stated above, in Q2 and Q3 of 2025, we saw two seven-digit ransoms across 29 total engagements. In the same quarters of 2026, we saw just one across 53 engagements. The mega-

paydays haven't disappeared, but they're becoming rarer in our dataset. What's replacing them is volume: more engagements, more victims, more pressure applied across a wider surface area. The big game hunt hasn't ended, but the hunting party has gotten much larger, and many of its members are content with smaller kills.

If the Money Is Down, Why Do the Numbers Keep Going Up?

If payment rates are declining and individual payment amounts aren't uniformly rising, then the per-incident expected value for attackers should be falling.

Basic economics would suggest that rational actors would move on to more profitable ventures. And yet, ransomware incident volumes are at all-time highs.

Several factors explain the paradox. First, the cost of conducting a ransomware attack has plummeted. Traditional ransomware is being replaced more and more by pure extortion models. Zscaler reported that data theft volumes surged 275% YoY, reaching nearly 900 terabytes of exfiltrated data. When your operational costs approach zero, even a lower expected payout per victim can generate meaningful returns at scale.

Second, ransomware has diversified beyond encryption. The shift toward pure data extortion – stealing sensitive information and threatening to publish it, without ever encrypting a file – means that attackers can run more campaigns simultaneously

with less infrastructure. The damage is no longer just the ransom. Rather, it's the operational disruption, the regulatory exposure, and the reputational fallout.

Third, and perhaps most importantly, the ecosystem is self-sustaining. The Ransomware-as-a-Service (RaaS) model means that even if top-tier operators see diminishing returns, the affiliate layer continues to churn. New groups emerge constantly because tools, playbooks, and infrastructure are readily available. The barrier to entry is low enough that even modest payouts justify the effort for operators in jurisdictions with limited law enforcement reach.

The Bottom Line

The “payments are down” narrative is comforting, but incomplete. Yes, more organizations are refusing to pay, and yes, that's a positive development driven by better backups, stronger incident response capabilities, and evolving regulatory pressure. However, ransomware operators aren't retreating. They're running more operations, targeting more victims, stealing more data, and relying on volume to compensate for lower per-incident conversion rates.

For executives and analysts evaluating their organization's risk posture, the message is straightforward: do not mistake a declining payment rate for a declining threat. The threat is growing, attackers are multiplying, and the cost of being unprepared has never been higher.

A Note on Our Data

A few caveats are worth noting. Our threat actor engagement data reflects cases where we were retained to conduct negotiations with ransomware groups, which inherently skews toward organizations with a higher propensity to consider payment, often because the situation is severe enough to warrant outside counsel and incident response. Our dataset is also limited to our engagements, which amounts to a window into a much larger ecosystem. Many of our clients carry cyber insurance, which can influence payment decisions in ways that differ from uninsured organizations. These factors don't invalidate the trends, but they do mean our numbers should be read as one data point in a larger picture, not as a definitive market survey.



Tinker, Model, Agent, Spy: How Ransomware and Cybercrime Groups Are Using AI in 2026

Our Baseline: AI as a Productivity Tool

Throughout each report this year, our assessments have encouraged a restrained view of the ways cybercriminals are adopting AI (artificial intelligence). In our [annual report](#) (released in January), we found that claims about “super-affiliates” deploying fully autonomous ransom-bots were overstated. We also noted that while experienced operators were incorporating AI/LLM (large language model) technologies into their work, usage remained rudimentary among less mature threat actors. Next, our [Q1 2026 report](#) (released in April) characterized AI largely as an attack surface rather than a threat itself. It reviewed a February 2026 supply chain attack on OpenClaw’s skills marketplace, noting that even the operator behind ShinySp1d3r described it as a rehash of Hellcat ransomware modified with AI tools. Then, our [Q2 2026 report](#) posited the idea that rather than creating novel capabilities, AI tools essentially just upskills actors, improving existing capabilities like those for reconnaissance, analysis of stolen data, drafting language for ransom negotiations, reducing language barriers, and increasing operational efficiency. While we stand by those previous assessments as relevant at their respective points in time, new information and new use cases have forced our Q3 assessment to adjust with the times.

Where We Are Seeing Change

AI is no longer just advising attackers. It is operating alongside them inside the intrusion. The clearest example of this shift was recently documented by [Gambit Security](#), as demonstrated by an affiliate of the Aurora ransomware group. Between April and May 2026, the affiliate reportedly used Cursor’s agentic coding assistant, running Anthropic’s Claude Sonnet, for hands-on intrusion work against at least ten victims. That work included reconnaissance, credential theft, and Kerberos/Active Directory (AD) escalation. To bypass the agent’s safety refusals, the affiliate repeatedly told it the activity was a “test.” The operator supplied stolen credentials for an existing route into the network, then directed Cursor to install and configure VPN or proxy tools, scan internal subnets, analyze AD privileges, attempt NTLM (New Technology LAN Manager) relay and certificate attacks, repair failed commands, and propose next steps. This marks a shift away from the chatbot-style use most of us know from daily work, where a person asks questions and acts on the answers, toward an agent that carries out intrusion steps with far less human involvement.

Opportunistic mass exploitation is now run by agents. In September, GreyNoise documented a likely Russian-speaking actor who used hundreds of AI agents, powered by [OpenAI’s Codex](#) harness and a DeepSeek model, along with publicly available offensive tooling, to compromise at least 440 PaperCut MF/NG instances belonging to 395 identified organizations in 48 countries. The actor went from an empty workspace to its first Remote Code Execution (RCE) achieved against a real victim in under four hours. It reached its first domain admin account two hours later, and once the full campaign launched, compromised at least 11 organizations in 26 seconds. GreyNoise could not determine whether the threat actor was developing access to hand off to affiliates or intended to pursue data theft or ransomware directly. For victims, that distinction matters less than the fact that access was produced at “machine speed.”

Anthropic’s data shows this operating model spreading across threat actor classes. [Anthropic’s September 2026 report](#) states that the autonomous operating model it first documented in November 2025, in a suspected state-sponsored campaign, has now spread to every class of actor it investigated. It adds that public offensive agent frameworks such as PentAGI reproduce much of the same scaffolding for anyone who downloads them. In suspected ShinyHunters affiliate activity, one operator breached a SaaS provider, extracted data belonging to

roughly 200 downstream customer organizations, and dumped over 2,100 Azure AD token sets spanning more than 40 corporate tenants in about 34 hours, with AI agents doing nearly all the work. Anthropic also notes that operators often may not understand each target environment, instead deferring to AI for the specifics. This takes the reduced barriers to entry we have long alluded to and raises the stakes. Operators can now direct intrusions they could not carry out themselves.

The first end-to-end agentic ransomware cases have appeared, albeit on a small scale. [In July, Sysdig identified JadePuffer](#) as what it believes is the first documented ransomware operation conducted entirely by an LLM agent. That agent handled reconnaissance, credential theft, lateral movement, persistence, privilege escalation, and encryption. The threat actor has since expanded, returning with ENCFORGE, a Go-based ransomware aimed at roughly 180 file extensions across the AI/ML (machine learning) stack, including model checkpoints, vector databases, training datasets, and embedding indices. Microsoft now tracks the threat actor behind these campaigns as [Storm-3168](#) and has observed a destructive stage in Azure that hit more than 100 storage accounts plus Key Vaults, Function Apps, VMs, and App Services in just seven minutes.

Timelines have compressed across the board. In September, [Google's Threat Intelligence Group \(GTIG\)](#) observed a case where threat actors compromised a cloud resource, then planned, built, and executed an agent-enabled mass credential harvesting campaign in under six hours. In the same vein, [CrowdStrike reports](#) that in 1H 2026, 88% of observed exploitation of vulnerabilities with a public Proof-of-Concept (PoC) exploit occurred within 48 hours of the PoC's release.

AI infrastructure has become “loot and target”. In their September threat report, Anthropic found that stolen AI credentials give operators three things at once: resale value, compute that runs at someone else's expense, and cover because activity is attributed to the legitimate key owner. ShinyHunters affiliates who obtained a victim's AI keys during an intrusion switched their own attack workloads onto those keys. CrowdStrike observed cloud-conscious eCrime activity surge by 171% as adversaries pursued credential theft, cryptomining, LLM abuse, and digital asset theft.

Where We Are Not Seeing Change

All said, the same old techniques are still in play. Anthropic's report details attacks involving stolen credentials, unpatched edge devices, exposed services, SQL injection, and phishing, with none dependent on unique techniques. Aurora's affiliate succeeded by exploiting the same AD weaknesses pentesters have flagged for years. JadePuffer's entry point was a known Langflow CVE, and analysts noted that none of its individual techniques were novel or complex. The model, however, was able to chain them into a complete attack against neglected, internet-exposed infrastructure.

Humans still pick targets and cash out. Humans have kept the decisions that matter most to them, including target selection, monetization, and review of results, with several of the most serious compromises coming from operations where a human directed every step. GTIG says it “has not yet observed threat actors deploying fully autonomous pipelines against targets in the wild.” In its six-hour case, AI accelerated activity after the victim environment had already been compromised, rather than independently discovering and exploiting an unknown vulnerability. Even in GreyNoise's PaperCut example, multi-day gaps occurred between initial access and domain admin because of inaction from the attacker.

AI speed has not produced uniform success. Of the 440 compromised PaperCut instances, the threat actor harvested credentials from 280, exfiltrated secrets from 137, and reached domain admin in only 12. Gambit's 30-50 percent speed estimate for Aurora came with the finding that most AI-issued commands still required iterative human refinement. Agents also introduce their own failure modes. The PaperCut actor hardcoded a 28-country exclusion list led by Russia, Belarus, and CIS states, and several agents ignored it and attacked victims from excluded countries anyway.

The Ransomware Business Model and Market Structure Remain Intact

AI has not displaced RaaS, leak sites, or double extortion, nor has it broken up the market’s concentration. In Q3 2026, the five most prolific groups claimed more than 35% of all recorded attacks, even as the number of active groups reached a record 112. The documented AI-heavy operators of Aurora sit well-outside that top tier. Our extortion-side findings still hold: AI’s most consistent ransomware-economy role remains helping threat actors process stolen data and pressure victims.

Reconciling the Two Views

Our reporting and historical assessments are based on a combination of open sources, public ransomware reporting, and GuidePoint’s incident response data. That provides us with unique insight weighted towards the tail end of successful intrusions. We most often see successful ransomware incidents and the ensuing response and negotiations from where we sit. Recent reporting from the vendors we’ve mentioned provide a different, unique, and invaluable perspective, formed from exquisite telemetry and collection exploited and analyzed by talented analysts. By integrating these perspectives and data from “left of bang,” we can and must update our picture of the threat landscape as it evolves.

Our assessment for Q3 2026 is that a view of AI in cybercrime as solely an “enabler, not [an] orchestrator” is no longer accurate as a blanket statement. AI *is* orchestrating the intrusion phase for a growing number of affiliate-tier and access-broker actors. It *is not* (yet) orchestrating the underlying business. Target selection, negotiation strategy, and monetization remain human-directed, and the largest ransomware brands show no public sign that AI has changed much of how they fundamentally operate.

Changing	Not Changing (Yet)
AI executing intrusion steps independently	Underlying TTPs (Tactics, Techniques, and Procedures) and initial access vectors
Accelerated agentic exploitation in minutes-to-hours	Human control of targeting and monetization
Increased use by low-skill operators; sophistication losing utility as an attribution signal	RaaS economics and top-tier group concentration
AI keys, agents, and ML assets as targets and loot	Effectiveness of hardening, patching, and identity controls (code is faster than inference)

WIIFD (What’s In It For Defenders)

This shift presents a new core threat: time. As threat actors increasingly employ AI and LLM tools that can process, infer, and act faster than humans, the window for error in detection and response narrows. Gone are the days in which an attacker may be stuck fiddling with a script or bashing their heads against a single control for hours or days at a time. We must depend on our detections, our controls, our code, to be faster than attacker-deployed inference.

While we are observing economic and scale shifts, these do not fundamentally alter the nature or manner of the threat; they merely increase volume. This poses outside threats to small and midsize businesses, non-profits, government, and other organizations with poor security resources far more than well-defended and well-resourced organizations in private industry. We expect to see higher loss-event frequency for organizations that were previously too small or too obscure to be worth a human operator’s time. Risk models that assume attacker effort is a limiting factor should be revisited.

The good news is that nothing in this reporting requires new categories of control. The Aurora affiliate got through because identity and certificate weaknesses were already present on the network. In PaperCut, one of three documented paths to domain admin was simply finding PaperCut running as a domain-admin service account and adding an attacker account to Domain Admins.

AI turns known weaknesses into faster losses, but it does not create new weaknesses on its own.

Defenders should resist two opposing analytic errors. One is dismissing AI-enabled threats because the techniques look familiar, which ignores the collapse in time-to-impact. The other is funding “AI-native” defenses ahead of brilliance in the basics: patch velocity, identity hygiene, and response automation, where AI-accelerated attacks are still won or lost.

How Defenders Can Respond

Short term (1 to 3 months)

- Treat internet-facing vulnerabilities with a public PoC or patch as a 48 to 72 hour problem, not a monthly cycle, starting with edge devices, print and file management servers, and self-hosted AI tooling such as Langflow.
- Audit Active Directory for the paths agents are using: ADCS (Active Directory Certificate Services) misconfigurations, NTLM relay exposure (enforce SMB and LDAP signing and Extended Protection for Authentication), noPac-class issues, and service accounts running with domain admin rights.
- Inventory every AI API key and agent integration, move them into a secrets manager, and extend secret scanning to mobile app builds, container images, and public repositories, since that is where the ShinyHunters affiliates were harvesting.

Mid-term (3 to 12 months)

- Close the gap between detection and containment. If an agent can go from access to domain admin in hours, containment steps such as token revocation, account disablement, and host isolation need pre-approved automation rather than a human approval chain.
- Build behavioral detections for “machine-speed” activity: bursts of enumeration and credential access across many hosts, rapid tool retries, and unexpected outbound LLM API traffic from servers that have no reason to make it.
- Bring AI assets under the same governance as production systems: sanctioned model access only, egress control for AI traffic, least-privilege permissions for agents and service principals, and vetting of third-party skills and MCP servers.
- Run tabletop exercises with compressed timelines, where initial access to domain compromise happens inside one working day and exfiltration happens before a SOC shift change.
- Update cyber risk quantification to reflect higher attack frequency against mid-market and previously low-profile environments.

Long term (12 months and beyond)

- Plan for increasing capability parity between criminal, hacktivist, and state actors, and build threat models around exposure and intent rather than sophistication.
- Move toward identity-first architecture: phishing-resistant MFA everywhere, tiered administration, and elimination of standing privilege, since identity was the decisive layer in almost every case above. Adopt agentic defense in the SOC deliberately, with the same governance you would demand of attacker-facing automation, because a human-speed response will not keep pace with agent-speed intrusion.
- Build secure-by-design requirements for internal AI deployments, including sandboxing, prompt-injection resistance, and credential isolation, before those deployments become the next Langflow.



Quarterly Wrap Up

Analyzing ransomware and cybercrime trends on a quarterly basis allows us the opportunity to spot emerging patterns, and Q3 2026 is no different in this regard.

We see diminished efficacy in at least some tactics: Clop's Windchill campaign reached fewer targets than MOVEit ever did, and the group was forced to contend with both a shrinking victim pool and inter-group rivalry with ShinyHunters. Clop's attack pattern over time depends heavily on zero-day exploits, but their approach to data extortion – "smash and grab" at scale – just does not seem to hit like it used to, especially when the data they grab is of fragmentary or low value.

As Defenders have internalized and responded to risks at their network perimeter – historically prime real estate for an attacker – we can see clear efforts to compensate by targeting the center instead. Specifically, SaaS platforms across every industry provide an opportunity to more easily export high volumes of data, with the added benefit of impacting multiple connected organizations (e.g., hundreds of schools through a centralized school district infrastructure). SaaS access is widespread across organizations and may be less rigidly administered, potentially allowing diligent social engineers the opportunity to skip privilege escalation and move straight to their actions on objectives, with readily available lateral movement afforded through authenticated credential theft. Defenders are forced to manage risk that lives in systems they don't own.

There's also a trend that may or may not bother you at first glance: fewer victims are paying, but the ecosystem is thriving anyway. Attacks are up double digits, threat actors are up double digits, and the cost of running an extortion campaign keeps falling toward zero as threat actors refine playbooks and employ AI. Ultimately, a lower conversion rate is poor consolation when the top of the funnel keeps getting wider.

What ties it together, and what we'll be watching into Q4, is time. Whether it's an AI agent going from access to domain admin in a single afternoon or a vendor compromise rippling out to a hundred downstream customers before anyone notices, the margin of error for a slow response to an intrusion keeps shrinking. Attackers are moving faster and casting their nets wider.

But there is good news. None of this requires fundamentally reinventing your security program. The controls that mattered in every case from this report – patch velocity, identity hygiene, response automation – are the same ones that mattered last year. Yes, the attackers got faster, but the fundamentals on which they depend didn't move. Organizations that get and stay serious about closing the gap between detection and containment will prove capable not only of outpacing a faster adversary but also outrunning other programs that may not have bothered to try. As with most technology changes, what we are seeing is incremental shifts and change – not yet a revolution.

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.



GUIDEPOINT®
SECURITY

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

R.GRITQ3.2610