NetDiligence®

Insights from the Cybersecurity Front Lines

NetDiligence® Ransomware Advisory Board: Q3 2026 Meeting - Executive Summary

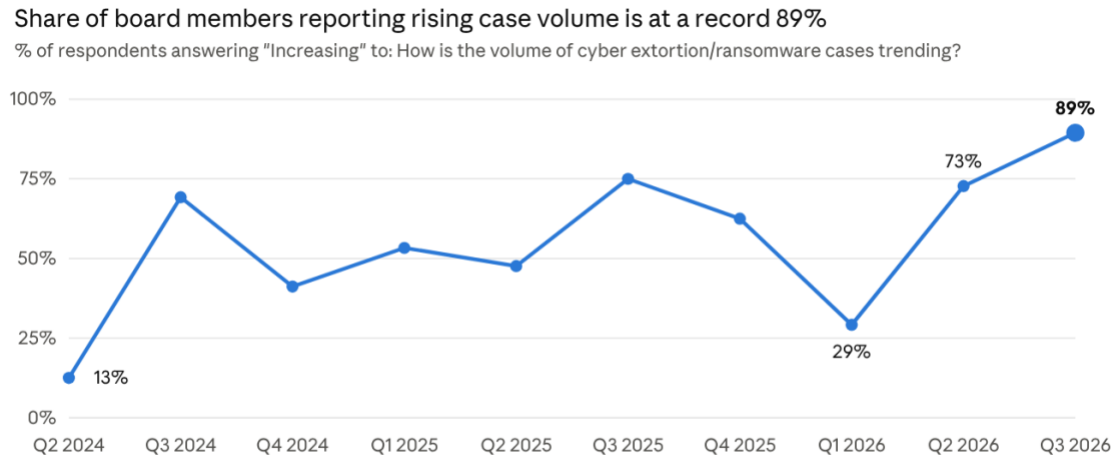
Each quarter, NetDiligence brings together incident response leaders for a private discussion on what they're seeing "in the trenches" across incident response engagements.

This briefing distills the most important Q3 2026 insights: how a "swarm" of new threat actors is driving case volume to its highest level in years, how AI has quietly become a standard part of the attacker toolkit, and why extortion is tipping toward data theft without encryption, with bigger opening demands and deeper discounts.

Prepared by Sherri Davidoff, Founder of LMG Security and Chair of the NetDiligence Ransomware Advisory Board.

Executive Summary

Case volume is surging. 89% of board members reported that cyber extortion and ransomware cases increased over the past three months, with 11% reporting flat volume and none reporting a decrease. This is the strongest "increasing" reading since the board began polling on volume in 2024, up from 73% last quarter and 29% in Q1 2026. Members corroborated the poll with their own data: leak-site postings are up roughly 20% year over year, public breach filings are up roughly 30%, and inbound incident calls at one large IR practice have risen by a similar margin.



Shown Above: NetDiligence RAB quarterly poll results, Q2 2024 - Q3 2026 (13-30 respondents per quarter)

The mix of attackers is changing. Previously unknown groups now account for roughly 30% of ransomware cases at one firm, not because established groups like Akira and Qilin are fading, but because a wave of new entrants has joined them. Members described the shift as "more snakes and fewer dragons," and pointed to a booming dark-web market for AI-enabled criminal tooling that lowers the barrier to entry.

© 2026 NetDiligence® All rights reserved. Unauthorized Disclosure Prohibited.

AI use by attackers has crossed a threshold. One year ago, every board member polled said they saw signs of attacker AI use "rarely." This quarter, 55% said they see it "often" or "almost always," and another 35% said "sometimes" (excluding simple translation and language polishing). Members stressed that AI is not creating novel attack paths; it is making known attacks faster, cheaper, and accessible to lower-skilled actors.

Extortion is tipping toward data-only. Only 16% of members said that a majority of their cases involved both encryption and a threat to publish, down from 58% last quarter and the lowest reading in the history of this poll. Opening demands are climbing into the millions, yet negotiated discounts of 60-70% are now routine, and negotiations are taking longer. Data-only extortion was reported as a majority of cases by 31% of members, up from 6% in Q3 2025.

Key Insights & Trends

1. A "swarm" of new threat actors is driving record volume Board members described a fragmented threat landscape in which unknown or newly observed groups make up roughly 30% of ransomware cases, alongside established players. Several members compared the moment to the splintering of major gangs after the 2022 invasion of Ukraine, but with a key difference: AI-enabled tools sold through dark-web forums and Telegram storefronts now let non-technical criminals buy deepfake kits, know-your-customer bypass services, jailbroken access to commercial chatbots, and ready-made attack tooling. Many of the resulting attacks are rudimentary, and some are more destructive than intended (for example, ransomware built incorrectly so that data cannot be decrypted). Members also reported "lone wolf" actors who, rather than demanding a ransom, offered to be hired as security consultants for around $10,000 to "help improve" the security of the victim they had just breached. Smaller, unaffiliated actors may also be deliberately staying below the law enforcement radar.

2. AI is now a routine part of the attacker toolkit Members agreed that AI has become "just Google" for attackers: the default place to turn for answers, code, and analysis at every stage of an attack. Concrete examples from recent cases included:

- Attackers feeding network reconnaissance output to an AI assistant and having it write custom data exfiltration scripts for non-standard data stores. Investigators found AI-generated comments in the code, including instructions like "if there's an error, upload it to me and I'll fix it." Members expect this to lead to larger data thefts from custom environments, not just standard exports from Salesforce or Snowflake.
- AI used to analyze stolen data and cross-reference it against the victim's industry, geography, and applicable regulations, in order to estimate the victim's cost of exposure and justify the ransom demand.
- Reported use of AI to rapidly build proof-of-concept exploits from firewall and VPN firmware, shrinking the window from CVE publication to working exploit from weeks to days or even hours.

Members emphasized that AI is not inventing new attack pathways. The techniques are the same well-documented ones responders have seen for years; they are simply executed much faster and by less skilled operators. One member offered a useful three-tier framing for clients: is AI *in* the attack (rare, a heavy footprint), *around* it (used to build tooling before launch, now common), or *controlling* it (an agent talking back to a model and adapting in real time, not yet observed)? Fully autonomous attacks

© 2026 NetDiligence® All rights reserved. Unauthorized Disclosure Prohibited.

were not reported this quarter. Members noted that autonomy is expensive ("tokenomics is a real thing") and predicted that theft of credentials to other organizations' AI systems will grow as attackers seek free compute.

Forensic artifacts of AI use are a new and important evidence category. Members are seeing AI-style comments and structure in attacker code, verbose local logs on hosts running AI tools (effectively a full chat transcript of prompts, responses, and tool calls), user-agent strings and source IP addresses tied directly to frontier AI labs, and attack timelines that are simply too fast for a human typing by hand.

3. Extortion is tipping toward data theft without encryption Multiple members reported that single-extortion cases (data theft with a threat to publish, but no encryption) now outnumber traditional double-extortion cases, with one large practice describing the ratio as roughly three to one. Active groups named in this category included ShinyHunters, Black File (now operating as Redact), and Helix. Many of these cases begin with compromise of a single endpoint or a stolen SaaS session token, after which the attacker takes everything that one user can reach and extorts the organization.

Even established encryption groups are experimenting. Qilin ran a five-to-six-week single-extortion campaign in which responders found no evidence of network intrusion at all, suggesting the group may have been monetizing old data; it then returned to encrypting victims in recent weeks. In one unusual case, Akira released a working decryptor to a small MSP victim after four weeks of hard-line negotiation and no payment. Members could not say whether this was a mistake or a decision to walk away.

4. Demands are higher, discounts are deeper, and negotiations take longer Members who do not normally see large cases reported that demands have jumped from the low hundreds of thousands to million and multi-million dollar figures. Negotiators confirmed that opening demands are substantially higher, but that discounts of 60-70% are now routine, and that the time from first contact to resolution has grown. Attackers are also using prior payments as leverage, telling victims that "other firms paid X" to anchor the negotiation. Members advised treating time as an ally and using data to calm stakeholders, especially the victim, while the number comes down.

5. Law firms and professional services are under sustained attack The legal and professional services sector was the most frequently cited target this quarter. Callback campaigns (a burst of spam email followed by a phone call from a fake "help desk" that walks the victim through installing a remote management tool) are increasingly aimed at law firms, and groups such as Luna Moth / Silent Ransom continue to target firms holding sensitive client data. Members noted that attackers appear to have learned that stealing client data from lawyers reliably gets paid.

The victims are strikingly consistent. Across members' cases, the socially engineered individual has almost always been a senior attorney, often in contracts or dispute practice, and typically over 60. These are professionals who are more likely to answer an unknown call and less likely to recognize the ruse. Law firm websites make the job easy: names, practice groups, clients, and direct phone numbers are all public.

6. Affiliates are shopping around, and new cartels are recruiting aggressively The Gentlemen, a splinter of Qilin, surpassed Qilin in attack volume in June 2026 after offering affiliates roughly 90% of ransom proceeds, far above the industry norm. The group has operationalized bring-your-own-vulnerable-driver (BYOVD) techniques, EDR killers, and rapid exploitation of CVEs. Its negotiators are inconsistent: in one

© 2026 NetDiligence® All rights reserved. Unauthorized Disclosure Prohibited.

case the operator was clearly U.S.-educated (fluent in American idiom, and well versed in the legal reasons counsel might decline to pay for data suppression), while in others the group behaved like "the Wild West," publishing data prematurely. Members also documented the same affiliate operating under multiple brands (DragonForce, Qilin, and others) with identical TTPs, including attempts to extort the same victim twice through different channels.

7. Evasion is moving into virtual machines, blockchains, and developer pipelines Members shared several technical trends that complicate detection:

- Qilin affiliates have been launching a hidden QEMU virtual machine on each compromised host to perform exfiltration and evade EDR; responders initially found no evidence on the machines until they discovered the VM. This technique was first seen about a year ago from a single affiliate and now appears to be spreading.
- Living off the land remains the norm; malware typically appears only at the end, when the encryptor is deployed.
- Command-and-control is becoming more durable. Members reported growing use of "EtherHiding," which stores C2 instructions on public blockchains rather than relying on DNS or IP-based infrastructure that can be blocked.
- Cybercriminals are adopting a tactic previously associated with nation-states: planting persistence in developer environments and infrastructure-as-code. Members reported heavy attacker interaction with GitHub instances and extortion tied to compromised developer environments. One global case involved backdoors inserted into cloud configuration templates and the destruction of the victim's AWS environment.

8. Exploitation is faster, and old vulnerabilities are back in play The window between vulnerability disclosure and exploitation continues to shrink; one recent vendor report measured time-to-exploit more than a month shorter than the prior period. Members cautioned that the window could eventually go negative as attackers use AI to find flaws before vendors do. At the same time, LockBit was observed scanning for 13 older CVEs dating back to 2021, a reminder that unpatched legacy flaws remain viable. Firewalls and VPNs remain the easiest targets, but members are seeing a broader range of exploited products, including web servers, analytics tools, and PBX systems. A common client request now is a retroactive 90-day hunt whenever a new exploit is published for a vulnerability that has been present in deployed software for a year or more.

Actionable Takeaways

1. Plan for data-only extortion and bigger opening numbers Update incident response and negotiation playbooks for cases with no encryption to recover from, where the entire leverage is stolen data. Prepare leadership for high opening demands and for a longer negotiation; the board's experience is that time and disciplined use of data drive deep discounts. Do not let an attacker's claim that "other firms paid X" anchor your decision.

2. Inoculate senior professionals against callback campaigns Law firms and professional services firms should brief partners and senior staff, not just junior employees, on the spam-then-phone-call pattern and the fake help desk script. Technically, block installation of unapproved remote management tools, allowlist the RMM products IT actually uses, and alert on new RMM agents appearing on endpoints. Review what the firm's own website reveals about who to call and what they work on.

© 2026 NetDiligence® All rights reserved. Unauthorized Disclosure Prohibited.

3. Add AI artifacts to your forensic checklist Incident responders should now routinely look for and preserve evidence of attacker AI use: AI-style comments and structure in scripts, local logs from AI coding tools (which can contain the attacker's full prompt history), user-agent strings and source addresses tied to AI providers, and timelines that are implausibly fast for a human operator. Expect larger data thefts from custom data stores, and scope exfiltration accordingly.

4. Compress your own patch window, and hunt backward Assume a working exploit will exist within days of a CVE for internet-facing firewalls, VPNs, web servers, and other edge devices. Patch those within days, not weeks, and do not deprioritize older CVEs that remain unpatched. When a new exploit is published for a flaw that has existed in deployed software for a long time, run a retroactive hunt (at least 90 days) for signs of earlier exploitation.

5. Treat developer environments and infrastructure-as-code as crown jewels Protect GitHub and other source repositories with phishing-resistant MFA and tight access controls. Review changes to Terraform, CloudFormation, and similar templates with the same rigor as production changes, and alert on unexpected modifications. Secure the CI/CD pipeline so that a planted backdoor cannot be pushed to production automatically.

6. Look for virtual machines and blockchain-based C2 Alert on unexpected hypervisor or QEMU processes and virtual disk files on workstations and servers, which may indicate a hidden VM performing exfiltration and EDR evasion. Ensure detection does not depend solely on blocking known C2 domains and IP addresses, since attackers are shifting to durable channels such as public blockchains.

7. Shrink detection and response time When AI lets an attacker move from access to exfiltration in minutes, a detection cycle measured in hours is too slow. Prioritize automated containment for high-confidence detections, and rehearse rapid session and credential revocation across SaaS and identity platforms.

Conclusion

The Q3 2026 discussion points to one overriding theme: speed. A broader, less skilled, AI-equipped population of attackers is generating record case volume, exploiting vulnerabilities within days of disclosure, and moving from initial access to large-scale data theft faster than human-paced defenses can respond. Extortion is shifting toward pure data theft, demands are opening higher, and negotiators are increasingly savvy about the legal pressures victims face.

The defensive priorities follow directly. Prepare for data-only extortion, protect the senior professionals and developer environments attackers are targeting, add AI artifacts to forensic playbooks, patch edge devices within days, and invest in automated containment that can act in minutes. The board will return to how AI is reshaping forensics and response at its Q4 meeting.

For questions about this summary or the Ransomware Advisory Board, contact NetDiligence at netdiligence.com/contact.

© 2026 NetDiligence® All rights reserved. Unauthorized Disclosure Prohibited.