RAPID7

REPORT

THE Q3 2025 THREAT LANDSCAPE

By Rapid7 Labs

CONTENTS

WELCOME TO THE Q3 2025 THREAT LANDSCAPE REPORT.

The past quarter saw relentless activity in the global cyber domain. From critical vulnerabilities exploited within days of disclosure, to the continued evolution of ransomware groups and nation-state actors, defenders faced a complex and shifting threat environment that demanded both speed and resilience.

Q3 was marked by a zero-day exploitation of Microsoft SharePoint servers and multiple critical vulnerabilities in Cisco products, both leveraged in rapid mass-exploitation campaigns. The race between disclosure and exploitation has never been tighter, with many organizations caught in the gap.

Beyond direct exploits, the Salesloft supply chain breach highlighted the growing fragility of trust in interconnected SaaS ecosystems. Attackers weaponized legitimate update and integration mechanisms, serving as a stark reminder that even trusted vendors can become an attack vector.

Ransomware groups like LockBit, Black Basta, and Qilin refined their extortion tactics, expanding their targeting toward critical services, manufacturing, and healthcare sectors. Notably, we observed an uptick in double- and triple-extortion methods, blending data theft, leaks, and sustained harassment of victims.

Geopolitically, nation-state activity such as APT29's ongoing credential theft campaigns and China-linked espionage targeting telecom and defense supply chains reinforced that cyber remains a central theater of global competition. Emergent threats like the Brickstorm espionage campaign further blurred the line between traditional intelligence gathering and cyber-enabled disruption.

Our analysis also underscores a persistent trend: Attackers increasingly target core business systems including ERP platforms, virtualization layers, and managed file transfer technologies. These are not random, opportunistic hits; they are calculated moves designed for operational disruption and leverage.

In this report, you'll find:

- Data-driven insights into top exploited vulnerabilities and their attacker ecosystems
- Trends in ransomware operations and victimology from our ransomware and dark web tracking
- Deep dives into notable APT campaigns and malware evolutions
- Guidance on defensive priorities drawn from our AttackerKB and Intelligence Hub correlations

As always, our mission remains to make the complex simple: Turn raw threat data into curated, actionable intelligence that empowers defenders to act decisively. Cyber threats continue to evolve, but so does our ability to anticipate, detect, and respond with precision. Thank you for joining us as we map the adversary's next move.

— Raj Samani, Chief Scientist, Rapid7

THE RANSOMWARE LANDSCAPE: TRENDS AND KEY PLAYERS

As we looked back on the first half of 2025 and ahead to Q3, we expected a prolonged powerscale rebalancing among threat actors, and this has largely come to pass. Many of the groups that at one time dominated suddenly went silent, while others announced various forms of retirement or similar pauses in activity. Other well-known groups such as Qilin and Akira have swooped in to fill the void; however, we've also seen the rise of "collectives" in which a few hardliners take the lead and many join in to participate in the action.

Our analysis of ransomware activity for Q3 2025 highlights the established dominance of threat actors, major groups forming alliances to further their goals of data exfiltration and extortion, and what may be a final farewell for long-absent ransomware groups. July through September was a time of power consolidation, and a broadening of scope for prolific threat actors such as Qilin, SafePay, and WorldLeaks.

While alliances can prove brittle and temporary, the previously mentioned team-up of Qilin and DragonForce with LockBit (sporting an all new version 5.0) should not be taken lightly. If the merging of tactics and infrastructure comes to fruition, it may well solidify Qilin and DragonForce's places in the top 10 not just for Q4, but 2025 as a whole.

At a glance

Looking at the top 10 ransomware groups in terms of number of leak site posts (Figure 1), we see that Qilin has retained the top position in Q3 2025, with Akira and INC Ransom quite a way behind in second and third place, respectively. Devman, WorldLeaks, and Everest are the only new additions to the top 10 this quarter.

In Q3, there were 88 groups actively posting to their leak sites, compared to 65 in Q2 and 76 in Q1. Of the 88 in Q3, 28 are new additions to the active group total. These new additions include (but are not limited to) Cephalus, Miga, Obscura, Radiant Group, and Yurei.

We also saw some ransomware groups go silent this past quarter. Of the 65 groups making leak posts in Q2, 19 are now inactive in terms of no visible leak posts for the months of July, August, and September.

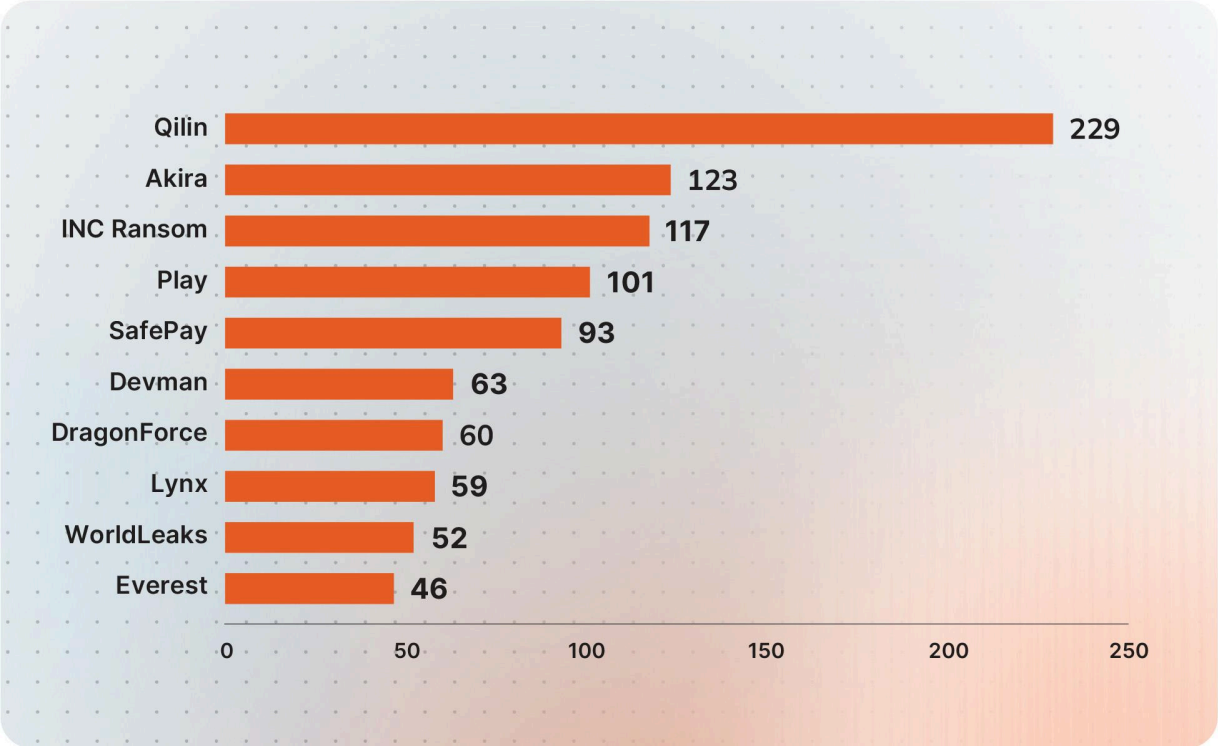

Figure 1: Top 10 ransomware groups by number of leak site posts, Q3 2025

Popular targets in Q3:

- Business services, manufacturing, healthcare, construction, and the financial sector were the most targeted industries in Q3 2025. The leader of the pack this quarter is business services, with 18% of posts containing these victims' data. Second place goes to manufacturing with 15% of posts, and healthcare sits third, with a total of 13%.
- Top regional targets include the United States with 67%, Germany in distant second with 6%, and the UK and Canada in joint third place with 5%. Italy (4%) and France (3%) fill the other two top spots.

Notable Trends

Rapid7 observed several notable trends at play in the ransomware landscape during Q3, including power consolidation among major players, evolving tactics in evading law enforcement, and innovations in operational models.

Hide and seek

The first half of 2025 saw new groups come and go, and many followed a trend of wanting to make a big, splashy announcement that they had arrived on the scene. Whether through slick branding or by making the most of a big compromise, being noticed by potential affiliates and other major players was the name of the game.

By contrast, Q3 brings a few tweaks to that approach. New groups aren't necessarily charging into battle via the loudest means imaginable. In fact, they're becoming more elusive where posting up their leaks are concerned. They're stripping out URLs and other identifiers such as victim names, and reducing everything down to information-lite screenshots instead. This is likely designed to throw law enforcement and researchers off the scent.

This game of cat and mouse also extended to major ransomware threat actors, even if not directly involved in said games, with time-wasting "rewards" designed to troll law enforcement agencies.

Sliding down the ladder

The first half of the year saw a visible climb to power among a handful of groups, while former major players were shut down, disbanded, or disrupted by law enforcement. As a result, threat actors such as Akira and Qilin have solidified their position at the top in Q3. Seemingly immovable objects such as Cl0p (or at least, someone claiming to be Cl0p), FunkSec, Cactus, and BianLian have all fallen into the mist.

Their radio silence has paved the way for domination in the ransomware space from Qilin and a handful of others. The unusual addition of non-SaaS SafePay into the mix may encourage other, smaller groups with a desire for success to also go down the fileless route with a focus on extortion over ransomware deployment.

Ones to watch

While ransomware groups remain fluid in their operations due to affiliate movement, global law enforcement activity, etc., the following are groups to keep an eye on as we finish out the year.

Qilin

As discussed above, the Qilin ransomware group has increased in both its visibility and impact as the year has progressed. The group has caused significant disruption in Q3, with attacks across multiple industries such as education, food/beverage, and healthcare. The cost to victims has ranged from exfiltration of large amounts of sensitive information (contracts, employee data, medical data), to disruption and shutdown of services in the most severe cases.

Qilin's double-extortion tactics, combined with a successful RaaS business model, have contributed to its success alongside other service additions and experiments such as so-called "legal advice," technical support, and negotiation assistance, where a senior member of the group works alongside a less experienced player as they move through the process of extorting the victim. Add to this the recent revelation that Qilin is strengthening its infrastructure and techniques by joining forces with LockBit and DragonForce (Figure 2), and we have the makings of a very strong close to the year if Qilin keeps up this kind of momentum.

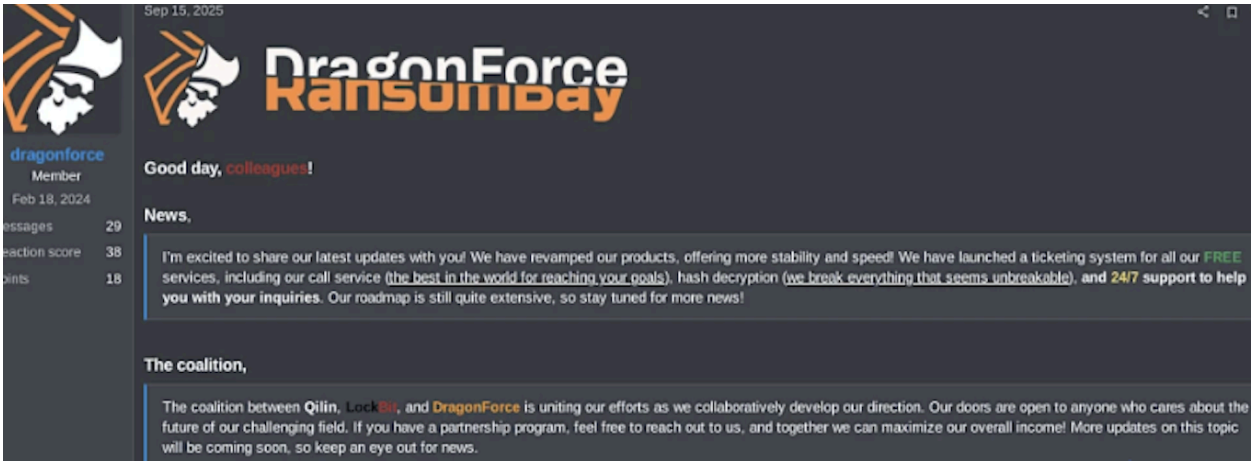

Figure 2: Screenshot of DragonForce admin post on the dark web site Ransombay, in which "the coalition between Qilin, LockBit, and DragonForce" is announced. The post includes the invitation: "If you have a partnership program, feel free to reach out to us, and together we can maximize our overall income!"

Crimson Collective

Crimson Collective, with a fondness for operating in cloud environments, first rose to prominence after claiming responsibility for the compromise of 28,000 private Red Hat repositories. This was followed in rapid succession by other activities targeting AWS instances, searching for leaked AWS credentials with the open source tool TruffleHog.

This group clearly has an eye for bigger things, and has also teamed up with the Scattered Lapsus$ Hunters collective, yet another example of threat actors pooling their talents and resources to keep the compromises — and the ill-gotten gains — flowing. Most recently, Crimson Collective is claiming to have breached Nintendo's internal systems, with alleged access to developer files, production assets, and backups. This collective has no intention of slowing down, and its fondness for seeking out exposed credentials in insecure code repositories likely means more compromises to come.

SafePay

First observed toward the end of 2024, SafePay became increasingly active in terms of threat posts during the first half of this year, and is now in fifth position of the Top 10 Ransomware Groups for Q3 (Figure 1). This is a solid showing against rivals such as INC Ransom and Play, especially as this group operates quite differently from most of the heavy hitters. SafePay does not make use of the RaaS model at all, preferring to keep all of their activities in-house and very much hands on. If you're compromised by SafePay, then you have direct experience of the threat actor and not "just" an affiliate.

This closed-off approach makes it more difficult for law enforcement to gain insight into the inner workings of the group. There is also little-to-no chance of the kind of inter-group duelling we've previously noted between DragonForce and RansomHub, given that SafePay seems to be very much about doing its own thing.

WorldLeaks

An interesting “new” group with quite a bit of backstory, WorldLeaks (first observed in January 2025) is an evolution of the threat actor Hunters International. Deciding that the world of ransomware, encryption, and double extortion was too risky, the group released free decryptors into the wild and focused on single extortion. That is to say, they narrowed their tactics to data exfiltration and the promise of releasing said data to the world at large.

This means no ransomware files, and no increased risk of being flagged by an organization’s intrusion detection defenses. WorldLeaks is more interested in gaining initial access than making a beeline for targeted data. Their leak site details company name, number of employees, revenue, and (notably) includes a view counter, no doubt to add a little more pressure onto the victim.

Nation-state activity: Espionage and exploitation

From the vantage point of Rapid7 Labs telemetry, Intelligence Hub correlations, and AttackerKB vulnerability tracking, nation-state operations in Q3 2025 were defined by a combination of stealth, supply-chain exploitation, and the tactical use of zero-day and near-zero-day vulnerabilities. The quarter demonstrated how established APT ecosystems continue to evolve their tradecraft — not necessarily by inventing new tools, but by refining their operational tempo, modularizing implants, and aligning technical intrusions with broader geopolitical objectives.

Russia

APT29 (Nobelium/Cozy Bear) remained one of the most active Russian-linked clusters, sustaining credential-theft campaigns that targeted Western diplomatic and defense organizations through OAuth token manipulation and misconfigured Azure app registrations. These operations reveal a maturing understanding of cloud identity systems, shifting away from spearphishing toward abusing trust relationships between SaaS and identity providers.

In multiple incidents observed also through partner telemetry, APT29 maintained access for months by embedding into legitimate synchronization processes, leaving minimal forensic trace and blurring the line between administrative behavior and compromise.

In a continuation of its campaign against Ukrainian critical infrastructure, the Russian state-sponsored group Sandworm deployed a new, highly destructive wiper malware known as PathWiper. The attack, which likely occurred around March but was analyzed in detail this quarter, targeted a Ukrainian critical infrastructure organization and is assessed with medium confidence to be the work of Sandworm, based on technical overlaps with previous wipers and the strategic nature of the target.

PathWiper represents a tactical evolution from previous Sandworm wipers, such as HermeticWiper. While HermeticWiper relied on exploiting a vulnerable legitimate driver to gain kernel-level privileges for direct disk access, PathWiper operates using standard Windows APIs. This suggests a shift in the deployment model. Instead of bringing their own privilege-escalation tool, the attackers focused on compromising an existing, high-privilege administrative platform. This method is stealthier and more efficient for mass deployment, as it leverages the target's own trusted infrastructure to deliver the destructive payload, highlighting an understanding of the victim's enterprise network management systems and capabilities.

Russian threat actors introduced a significant evolution in malware capabilities with the deployment of LAMEHUG, a tool used by a Pawn Storm-affiliated group that incorporates a large language model (LLM) to dynamically generate commands. This represents a tactical leap, creating more evasive and unpredictable malware.

China

China-nexus actors such as APT41 (Wicked Panda), and Volt Typhoon broadened their targeting scope beyond traditional espionage into strategic infrastructure domains. APT41 was observed leveraging code-signing abuse and living-off-the-land binaries to maintain persistence within software vendors that supply government and energy clients. This suggests a continued emphasis on third-party infiltration, over direct compromise.

Volt Typhoon, on the other hand, persisted in targeting maritime logistics and regional telecom operators, deploying custom router implants and SOHO exploitation chains that emphasize pre-positioning rather than disruption. Several Volt Typhoon operations aligned closely with regional military exercises in the Indo-Pacific, reinforcing how cyber operations are being synchronized with real-world geopolitical events.

One of the most notable developments this quarter was the emergence of Brickstorm, a newly detailed espionage campaign attributed to an Iranian nexus actor set. As analyzed by Google's Threat Intelligence Group and confirmed by independent telemetry, Brickstorm demonstrated highly tailored phishing with lure documents impersonating defense contractors, followed by deployment of a modular PowerShell-based loader designed to evade endpoint detection and response (EDR) visibility.

What makes Brickstorm noteworthy is its flexible command architecture; its operators leveraged compromised public infrastructure and cloud hosting accounts for staging, rapidly rotating servers to avoid static attribution. This adaptive approach mirrors techniques once characteristic of Russian APT tooling, suggesting cross-pollination of tactics among state-sponsored ecosystems.

Prior to the public disclosure of the F5 breach, the analysts at Google documented UNC5221 compromising network-edge appliances and servers — often Linux or BSD-based systems where traditional EDR agents are unavailable — deploying the BRICKSTORM backdoor to gain long-term persistence.

In the publicly disclosed F5 incident, F5 confirmed that an unauthorized actor gained access to its internal development and engineering knowledge-management systems, exfiltrating files from the BIG-IP product development environment beginning on or around August 9, 2025. Reports link this breach to BRICKSTORM/UNC5221, stating that the attackers used custom malware (BRICKSTORM) and a long dwell time (in some cases, more than 12 months) to harvest source code, internal vulnerability information, and configuration artifacts.

DPRK

The Kimsuky and Lazarus groups maintained consistent financial and espionage campaigns targeting cryptocurrency infrastructure and defense supply chains. Lazarus' adoption of artificial intelligence (AI)-generated deepfake personas for social engineering and Python-based loaders signed with stolen developer certificates marked a concerning step forward in operational deception.

Throughout Q3 2025, the DPRK-affiliated group Void Dokkaebi (also tracked as Famous Chollima and CL-STA-0240) continued its financially motivated "Contagious Interview" campaign. This operation represents a strategic focus on compromising the software supply chain by directly targeting developers in the technology and cryptocurrency sectors. Rather than attacking end-user applications or network infrastructure, this campaign poisons the well of the open-source ecosystem, exploiting the inherent trust developers place in public code repositories.

This approach is highly scalable and difficult to defend against, as it leverages social engineering to turn a developer's own tools and workflows into an attack vector. By compromising a developer's machine or tricking them into using a malicious software package, the threat actor can achieve its immediate goal of stealing cryptocurrency or credentials. More critically, this access creates the potential for a much broader downstream supply chain attack, where malicious code could be injected into a legitimate software project and distributed to thousands of unsuspecting users.

Building on our ongoing tracking of DPRK-linked cyber operations, Rapid7 Labs conducted an in-depth investigation into the network of DPRK-affiliated IT workers active across freelance and outsourcing platforms. By correlating personal data, infrastructure reuse, and shared social graph patterns, we identified clusters of interlinked profiles that reveal a coordinated workforce strategy designed to generate revenue and enable broader intrusion operations.

The analysis uncovered overlapping identifiers — including reused email domains, cryptocurrency wallets, and code repository handles — connecting individuals advertising themselves as developers or security consultants to known Lazarus and Kimsuky infrastructure.

Figure 3 illustrates these hidden interconnections, exposing how front-company accounts and contractor profiles form a distributed operational ecosystem that supports both financial and espionage objectives.

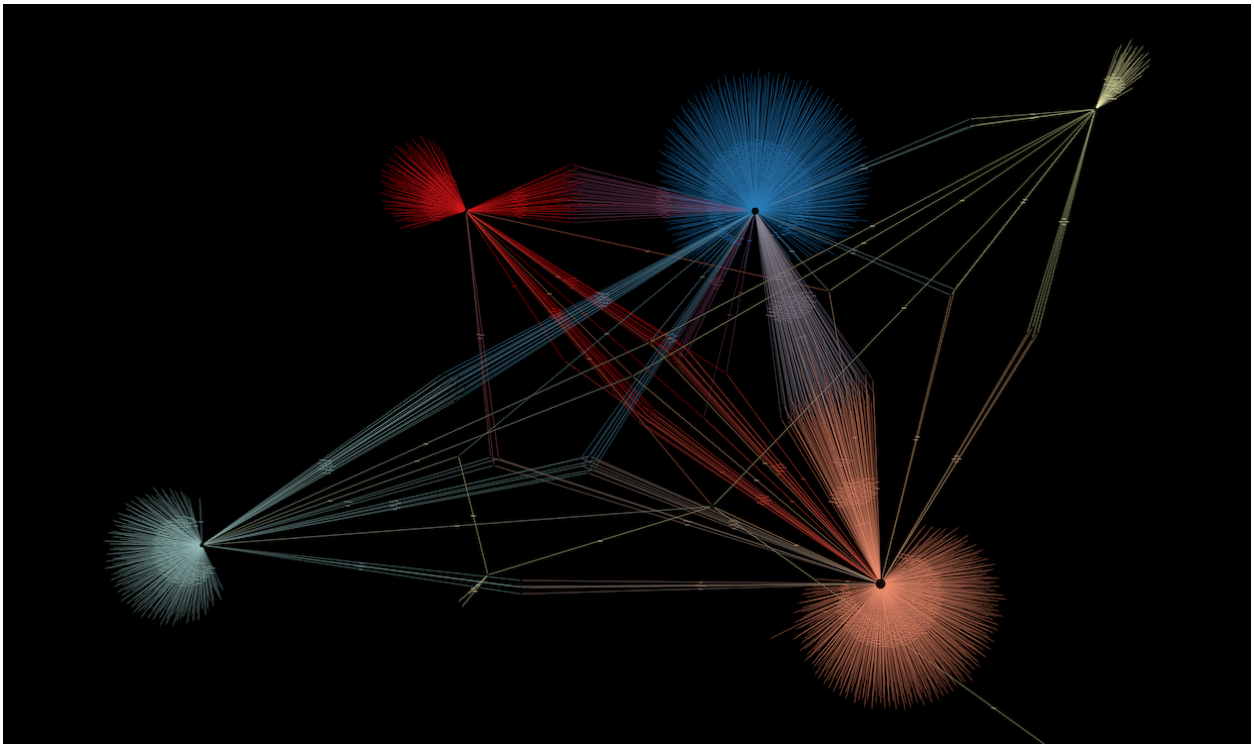

Figure 3: Each color cluster represents the Github account of a DPRK IT worker, where correlations denote shared projects, code, friends, etc.

Across these campaigns, Rapid7 Labs observed a growing interplay between exploitation and espionage, with APT actors rapidly operationalizing newly disclosed vulnerabilities, often before public proof-of-concept code appeared. The correlation between AttackerKB's "exploited-in-the-wild" tags and Intelligence Hub nation-state activity spikes suggests that vulnerability weaponization timelines are compressing, particularly in high-value network-edge and virtualization technologies.

Collectively, these campaigns illustrate a broader strategic evolution: Nation-state actors are moving from loud, high-impact compromises toward quiet infiltration of trust chains, identity systems, and supply-chain dependencies. Their objective is not always immediate disruption, but strategic persistence — remaining embedded and invisible until political calculus demands activation.

CYBERSECURITY INCIDENTS: IMPACTS AND KEY MITRE ATT&CK TECHNIQUES

The third quarter of 2025 saw a diverse range of sophisticated cyber incidents, primarily characterized by compromised credentials, extensive lateral movement, and the pervasive use of ransomware. Threat actors consistently leveraged common vulnerabilities and misconfigurations, particularly in remote access services, to gain initial access and establish persistence within victim networks.

The Rapid7 Incident Response (IR) team observed two examples of Crimson Collective activity in September, with attacks focused on data exfiltration and extortion of victims. These attacks made use of the popular open source tool TruffleHog, designed to detect insecure credentials stored in GitHub repositories. After initial access was gained, the threat actor would create new users and escalate privileges, attaching policies to the accounts which would grant full access to AWS services and resources. The end result of these network intrusions was a ransom note, and a warning to "tell your superiors."

Elsewhere, we observed an uptick in intrusions following a campaign used to gain initial access via SonicWall devices. Our investigations suggested that the Akira group was potentially utilizing a combination of three security risks — which can include the default configuration of Virtual Office Portal, and certain configurations of the SSLVPN Default Users Group — to gain unauthorized access and conduct ransomware operations.

At a glance

The incidents in Q3 2025 underscore the critical need for robust cybersecurity defenses. The impact on affected organizations ranged from significant operational disruption due to ransomware to potential data exfiltration and reputational damage.

- **Remote access exploitation and credential compromise:** A recurring theme across multiple incidents was the exploitation of remote access services, such as VPNs (SonicWall SSLVPN, Cisco ASA VPN, FortiGate firewall VPN), and remote desktop protocol (RDP). In many cases, these services lacked multi-factor authentication (MFA), making them vulnerable to brute-force attacks and the use of compromised or weak credentials. The social engineering of IT service desks to reset credentials and remove MFA also highlights a significant human element vulnerability.
- **Ransomware dominance:** Ransomware, particularly variants associated with the Akira RaaS (ransomware-as-a-service) group and the INC Ransom group, continued to be a primary objective for threat actors. Multiple incidents in Q3 demonstrated the significant impact of these attacks, often involving data encryption and deletion of backups.

- **Advanced tactics**: Several incidents showcased tactics, techniques, and procedures (TTPs) indicative of more advanced adversaries. This included the deployment of web shells and the extraction of data, the use of sophisticated tools like Impacket for remote execution and credential dumping, and the abuse of legitimate tools for malicious purposes (e.g., WinRAR for data staging, Rclone for data exfiltration, SoftPerfect Network Scanner for network discovery, Cloudflared for tunneling, and AnyDesk for persistence). Anti-forensic techniques like timestomping were also observed.
- **Internal reconnaissance and lateral movement:** Once initial access was gained, threat actors consistently performed extensive internal reconnaissance, utilizing tools like Advanced IP Scanner, Netscan, and Active Directory enumeration tools (ADExplorer). Lateral movement was predominantly achieved through RDP, often with compromised privileged credentials, to access sensitive resources such as domain controllers, file servers, and backup systems.
- **Shared threat actor targets:** The most impacted sectors in Rapid7-investigated incidents for Q3 were retail, real estate, and manufacturing, with industries impacted overall aligning with ransomware trends (e.g., construction, legal services, healthcare, and other industries found in the ransomware crosshairs).

Initial access

Initial access vectors frequently observed by the Rapid7 Incident Response team in Q3 included:

- **Exploitation of remote access services:** This was a primary vector, particularly through vulnerable VPNs (SonicWall SSLVPN, Cisco ASA VPN, FortiGate firewall VPN) and remote desktop protocol (RDP). The absence of multi-factor authentication (MFA) on these services significantly contributed to their compromise.
- **Compromised credentials:** Threat actors gained initial access by using stolen or weak credentials, often obtained through brute-force attacks or other means.
- **Social engineering:** In some cases, IT service desks were socially engineered to reset credentials or remove MFA, allowing attackers to bypass security controls.
- **Web shell deployment:** The deployment of web shells on compromised servers was also observed as an initial access method, providing persistent access and control.

Top 10 MITRE ATT&CK techniques for Q3

Monitoring the top MITRE ATT&CK techniques within a quarter is helpful in understanding the evolving threat landscape and prioritizing defensive strategies. By analyzing the prevalence of these techniques, organizations can gain actionable insights to strengthen their security posture, allocate resources effectively, and proactively defend against common attack methodologies.

More specifically, this knowledge can inform the development of specific detection rules, enhance incident response playbooks, and guide security awareness training to address the most frequently observed TTPs.

1. **T1078 - Valid accounts:** This is evident through the widespread use of compromised credentials and the exploitation of remote access services lacking MFA.
2. **T1133 - External remote services:** The exploitation of VPNs (SonicWall SSLVPN, Cisco ASA VPN, FortiGate firewall VPN) and RDP for initial access clearly falls under this technique.
3. **T1059 - Command and scripting interpreter:** The use of tools like Impacket for remote execution and credential dumping, as well as the deployment of web shells, indicates the use of command and scripting interpreters.
4. **T1003 - OS credential dumping:** Mention of extracting MachineKey data and credential dumping (e.g., via Impacket) directly points to this technique.
5. **T1021 - Remote services:** Lateral movement often occurred through RDP, which is a form of remote services.
6. **T1047 - Windows management instrumentation (WMI):** While not explicitly stated, the use of tools for remote execution and system interaction often leverages WMI.
7. **T1087 - Account discovery:** Active Directory enumeration tools (ADExplorer) were used for internal reconnaissance, which is a form of account discovery.
8. **T1018 - Remote system discovery:** Tools like Advanced IP Scanner and Netscan were used for network discovery, falling under remote system discovery.
9. **T1560 - Archive collected data:** The use of WinRAR for data staging and Rclone for data exfiltration suggests the archiving and compression of data before exfiltration.
10. **T1566 - Phishing (and other forms of social engineering):** The social engineering of IT service desks to reset credentials and remove MFA is a clear example of phishing or other social engineering tactics.

VULNERABILITY INTELLIGENCE: EMERGENT THREATS AND EXPLOITATION TRENDS

Looking at what vulnerabilities were reported in Q3 as having been exploited in the wild for the first time, we can see this figure is starting to trend down on previous quarters. Utilizing our community-based vulnerability intelligence platform, AttackerKB, Rapid7 Labs tracked a total of 53 vulnerabilities, all of which were first reported as exploited in the wild in Q3. This represents a notable drop on previous quarters, ultimately trending down quarter on quarter from this time last year (Figure 4).

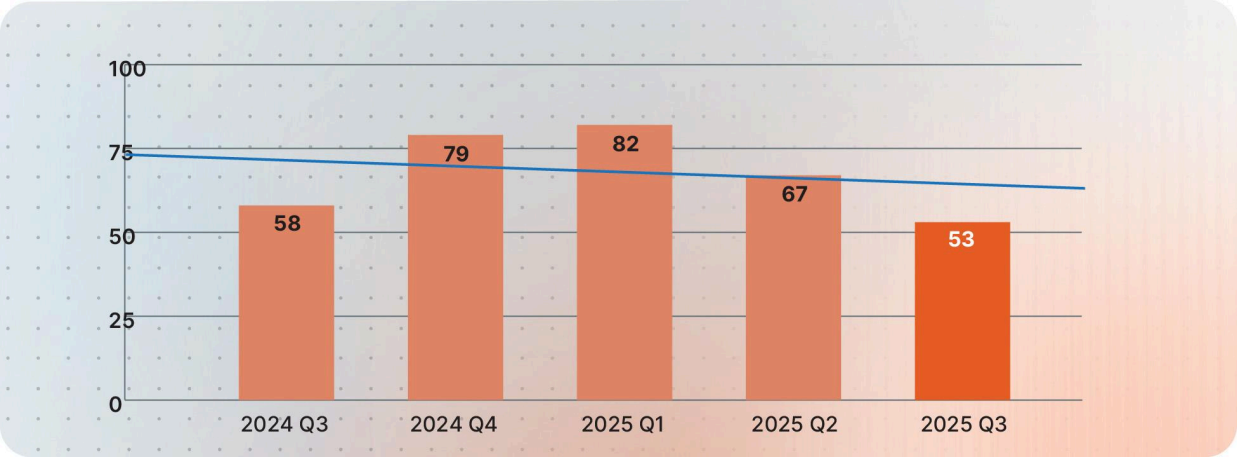

Figure 4: Number of vulnerabilities reported to be exploited in the wild for the first time, by quarter

While this finding is encouraging at a macro level, drilling into this data shows some troubling patterns that organizations need to consider in their strategy for vulnerability management (VM) and remediation.

If we examine the 53 vulnerabilities, we can see the delta, measured in days, from the initial disclosure of each vulnerability's CVE record to the date of its first reported in-the-wild exploitation. It's noteworthy that while the majority of these vulnerabilities have this delta in and around 0, i.e. newly discovered vulnerabilities exploited either as zero day or shortly after disclosure, a significant number of outliers exist. The delta of those outliers is measured in several thousand days, i.e. significantly older known vulnerabilities disclosed in some cases over a decade ago (Figure 5).

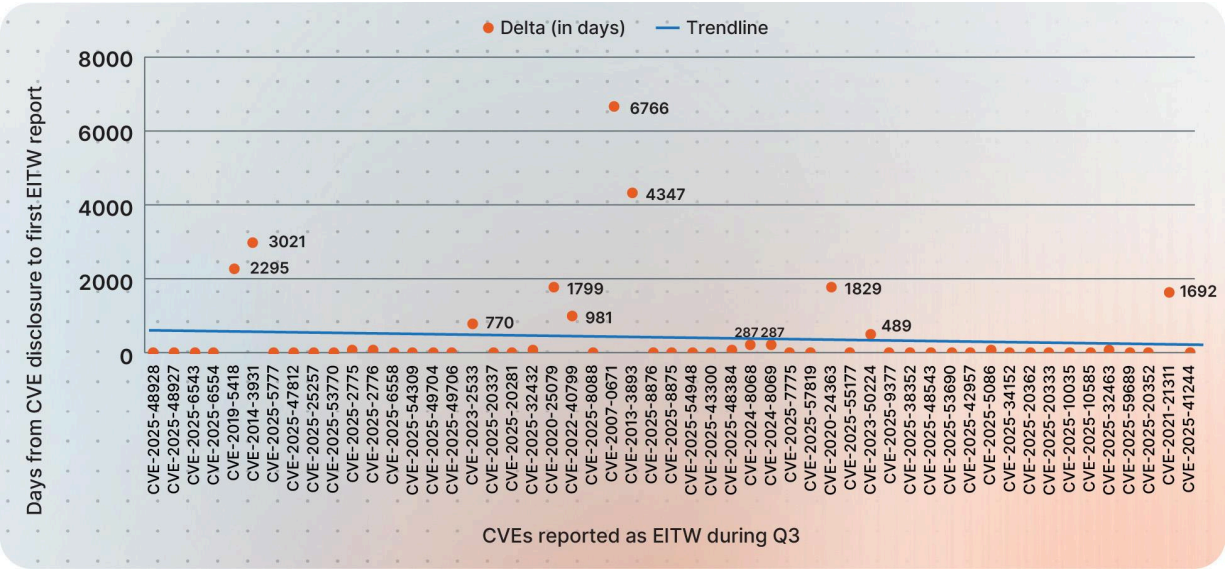

Figure 5: Delta of the CVE disclosure date to the earliest exploited-in-the-wild date, Q3 2025

CVE-2007-0671, for example, is 18 years old and received its first known exploitation-in-the-wild report in this quarter, via CISA's KEV list. In fact, 10 vulnerabilities from Q3 were disclosed more than a year ago, before being officially reported as exploited in-the-wild this quarter.

This highlights several issues. Firstly, relying solely on lists of known exploited vulnerabilities to action remediation plans is useful, but brittle. Older vulnerabilities that are known to be exploited and predate the existence of these lists may not necessarily be on these lists. CVE-2007-0671 — only added to CISAs KEV list this quarter — was actually noted as being exploited in-the-wild by the vendor circa 2007. Secondly, this highlights how remediation plans must not ignore patching historical vulnerabilities, which remain as potent attack vectors for threat actors even many years later.

Taking a look at the top 10 Common Weakness Enumerations (CWE) that have been used to describe the root cause of the 53 vulnerabilities (Figure 6), we see that CWE-502, aka unsafe deserialization, has a notable lead when compared to the other entries. This vulnerability class is found in memory safe languages like Java and C#, both very common within enterprise software products. They can provide an attacker with a reliable mechanism to achieve arbitrary code execution through the unsafe deserialization of malicious deserialization gadgets. Attackers favor this over vulnerability classes such as memory corruption, due to the reliability of exploitation.

CWE-78 and CWE-77 are in second and third place, both of which correspond to similar forms of command injection. Again, this is a logic level vulnerability class giving an attacker a reliable mechanism to achieve arbitrary code execution.

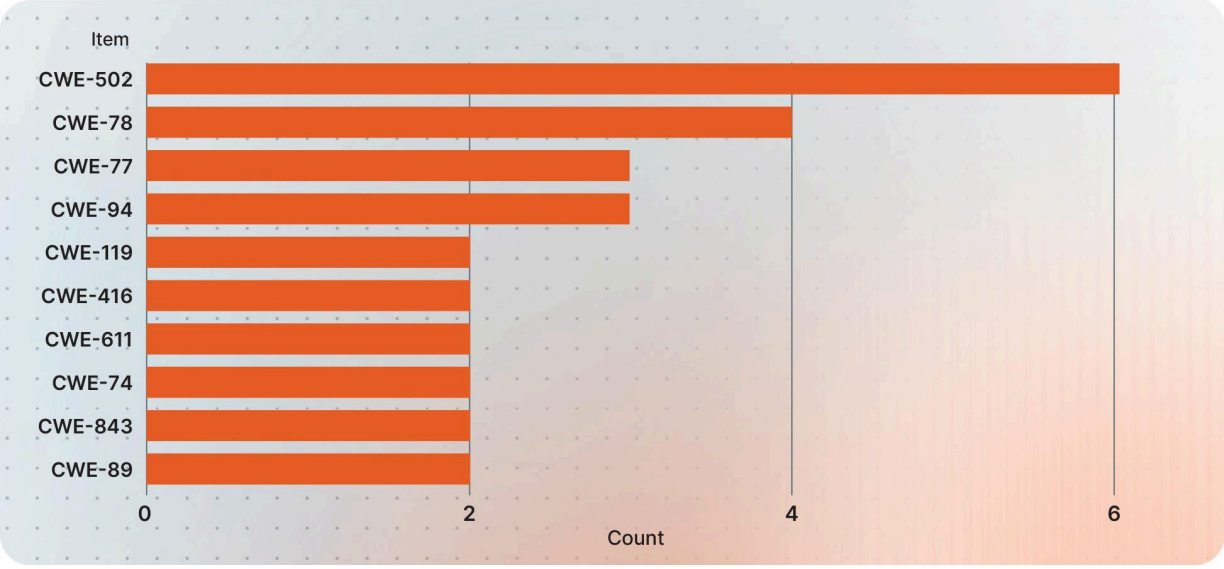

Figure 6: Top 10 CWEs for reported exploited in-the-wild vulnerabilities, Q3 2025

While the amount of newly reported exploited in-the-wild vulnerabilities may be trending down, Q3 saw a number of significant threat actor campaigns exploiting software vulnerabilities as part of Rapid7's Emergent Threat Response (ETR) program. Rapid7's ETR program aims to provide a rapid and holistic response to high profile software vulnerabilities that are either currently, or are likely to be soon, broadly exploited in-the-wild. Our ETR program in Q3 saw five such threats, all of which had been exploited by threat actors as zero-day vulnerabilities, i.e. before the availability of vendor patches.

- **CVE-2025-54309 - CrushFTP zero day exploited in the wild:** This popular enterprise file transfer product was exploited via a zero-day vulnerability, which was subsequently disclosed by the vendor in early July, 2025. This product has seen multiple vulnerabilities exploited in the wild over the past several years. File transfer solutions are a popular product group for attackers to target, not just as an initial access vector, but for data theft and extortion by ransomware groups.

- **CVE-2025-53770 - Zero-day exploitation in the wild of Microsoft SharePoint servers:** An exploit chain, dubbed ToolPane, comprised of CVE-2025-49706 and CVE-2025-49704, was initially developed and demonstrated at the 2025 Pwn2Own Berlin zero-day exploit competition. The exploit chain was leaked to a threat actor via unknown means. The threat actor began to exploit the vulnerabilities shortly before the regular Microsoft July 2025 patch release.

  Two follow-on vulnerabilities, CVE-2025-53770 and CVE-2025-53771, which were patch bypasses of the original exploit chain, were subsequently patched via an out-of-band patch release by Microsoft. Due to the pervasive deployment of SharePoint across multiple industry verticals, this threat campaign was likely the most significant and highly impactful incident in Q3.

- **CVE-2025-7775 - Critical NetScaler vulnerability exploited in the wild:** Citrix NetScaler once again saw in-the-wild exploitation of a new zero-day vulnerability by a threat actor performing targeted attacks. Exploitation in this case was highly targeted, and no public exploit code has yet been published online, preventing broad exploitation from multiple threat actors occurring.

- **CVE-2025-10035 - Critical unauthenticated RCE in GoAnywhere MFT:** Another popular enterprise file transfer solution was exploited with a zero-day vulnerability. In this case, exploitation has been attributed to the Storm-1175 threat actor, and at least one instance of the deployment of the Medusa ransomware was reported. Rapid7 Labs published a detailed root cause analysis of this vulnerability and discovered it was a complex exploit chain of three separate issues, along with an unusual requirement that an attacker know a specific private key, something the vendor has not yet clarified.

- **CVE-2025-20333, CVE-2025-20362, CVE-2025-20363 - Multiple critical vulnerabilities affecting Cisco products:** In late September, 2025, Cisco published patches for multiple critical vulnerabilities affecting their Adaptive Security Appliance (ASA) and Firewall Threat Defense (FTD) lines of products. Two of these vulnerabilities, CVE-2025-20362 and CVE-2025-20333, comprise a complex exploit chain that has been exploited in-the-wild by an as-yet-unknown threat actor in what appears to be a highly targeted attack. Due to the enormous deployment of ASA appliances worldwide, this threat poses a significant risk for organizations who have not yet patched. Rapid7 Labs published a detailed root cause analysis of this vulnerability and discovered that CVE-2025-20362 is trivial to exploit and is a patch bypass of an older vulnerability, CVE-2018-0296. Rapid7 Labs assessed that CVE-2025-20333 is a complex memory corruption vulnerability, and exploitation will be non-trivial.

The threat activity from Q3 highlights how zero-day vulnerabilities are continuing to be used by threat actors in targeted campaigns. Defending against zero days is inherently complex, as by definition vendor supplied patches are not available.

Organizations must leverage best practices to ensure their environments are hardened against these attacks. This includes following vendor guides to harden these environments, limiting what is exposed and removing features, services, and appliances from the internet perimeter to reduce your edge-exposed attack surface as much as possible. Continuing to monitor and threat hunt in your environments to detect anomalous activity is also essential.

Q3 also highlighted how older vulnerabilities are being exploited, often for the first time. This highlights the need for VM and the importance of remediation of known CVEs across your environment.

Ultimately, a balance must be found to effectively handle this complex and often changing threat landscape. Vulnerability intelligence can help close the gap between the known and unknown threats and their impact.

AI-SUPPORTED THREATS: SOCIAL ENGINEERING AND EVASIVE MALWARE

In Q3 we saw continued evidence of how AI is reshaping both the tempo and tactics of cyberattacks. Threat actors are no longer experimenting — they are operationalizing AI to enhance deception, automate intrusion workflows, and evade detection at scale. From leveraging machine learning (ML) improvements and increasingly convincing deepfakes for social engineering to dynamically-generated malware code, threat actors are benefitting from AI's ability to dramatically decrease the barrier to conducting sophisticated attacks.

For defenders, what was once a theoretical risk is now a practical, recurring reality. This underscores that AI is not just augmenting existing threats, it is redefining them. Defenders must now assume that adversaries can synthesize, adapt, and disguise faster than traditional detection or awareness training can keep pace.

AI-driven social engineering

Threat actors are leveraging generative AI to automate the creation of highly convincing phishing lures, deepfake audio and video for vishing campaigns, and tailored content for influence operations. This escalates the threat beyond what traditional security awareness training can address. Defense must become multi-layered, combining technical controls like phishing-resistant MFA(e.g., FIDO2) with continuous training that specifically educates employees on how to identify AI-generated content as well as sophisticated social engineering tactics.

AI-generated and evasive malware

The emergence of malware like Pawn Storm's LAMEHUG, which uses an LLM to dynamically generate its command and control logic, signals a new frontier in malware development. Such techniques are designed to evade signature-based and static analysis detection engines. Consequently, defensive postures must pivot from a reliance on known indicators of compromise (IOCs) to a focus on behavioral analysis, EDR, and network monitoring.

STRENGTHENING SECURITY: KEY RECOMMENDATIONS

Based on the analysis within this report, the following recommendations are crucial for strengthening organizations' security postures and enhancing their ability to prevent, detect, and respond to sophisticated cyber threats.

- **Develop comprehensive ransomware preparedness strategies**, including:
  - Implementing immutable backups, regularly testing recovery procedures, and ensuring backups are isolated from the production network.
  - Deploying and optimizing EDR solutions to detect and respond to ransomware activities early in the kill chain.
  - Segmenting networks to limit lateral movement in the event of a breach.
  - Developing and regularly testing incident response plans, including communication strategies for various stakeholders.

- **Prioritize the implementation and strict enforcement of MFA** across all critical systems, remote access points, and privileged accounts. This includes phishing-resistant technical controls as well as robust processes for credential management, regular audits of access privileges, and secure procedures for IT service desk operations to prevent social engineering attacks.

- **Maintain a mature VM program** with continuous scanning and timely patching of all public-facing applications and network devices. Focus on hardening configurations to minimize attack surfaces, and ensure that historical vulnerabilities are also addressed, as they remain potent attack vectors.

- **Leverage best practices to harden environments against zero-day attacks.** This includes following vendor guides, limiting exposed services and features, and continuously monitoring and threat hunting for anomalous activity.

- **Provide continuous and engaging security awareness training** for all employees, focusing on both traditional and AI-driven social engineering tactics, and the importance of reporting suspicious activities.

About Rapid7

Rapid7 is creating a more secure digital future for all by helping organizations strengthen their security programs in the face of accelerating digital transformation. Our portfolio of best-in-class solutions empowers security professionals to manage risk and eliminate threats across the entire threat landscape from apps to the cloud to traditional infrastructure to the dark web. We foster open source communities and cutting-edge research–using these insights to optimize our products and arm the global security community with the latest in attacker methodology. Trusted by more than 11,000 customers worldwide, our industry-leading solutions and services help businesses stay ahead of attackers, ahead of the competition, and future-ready for what's next.