VulnCheck

State of Exploitation

1H 2026

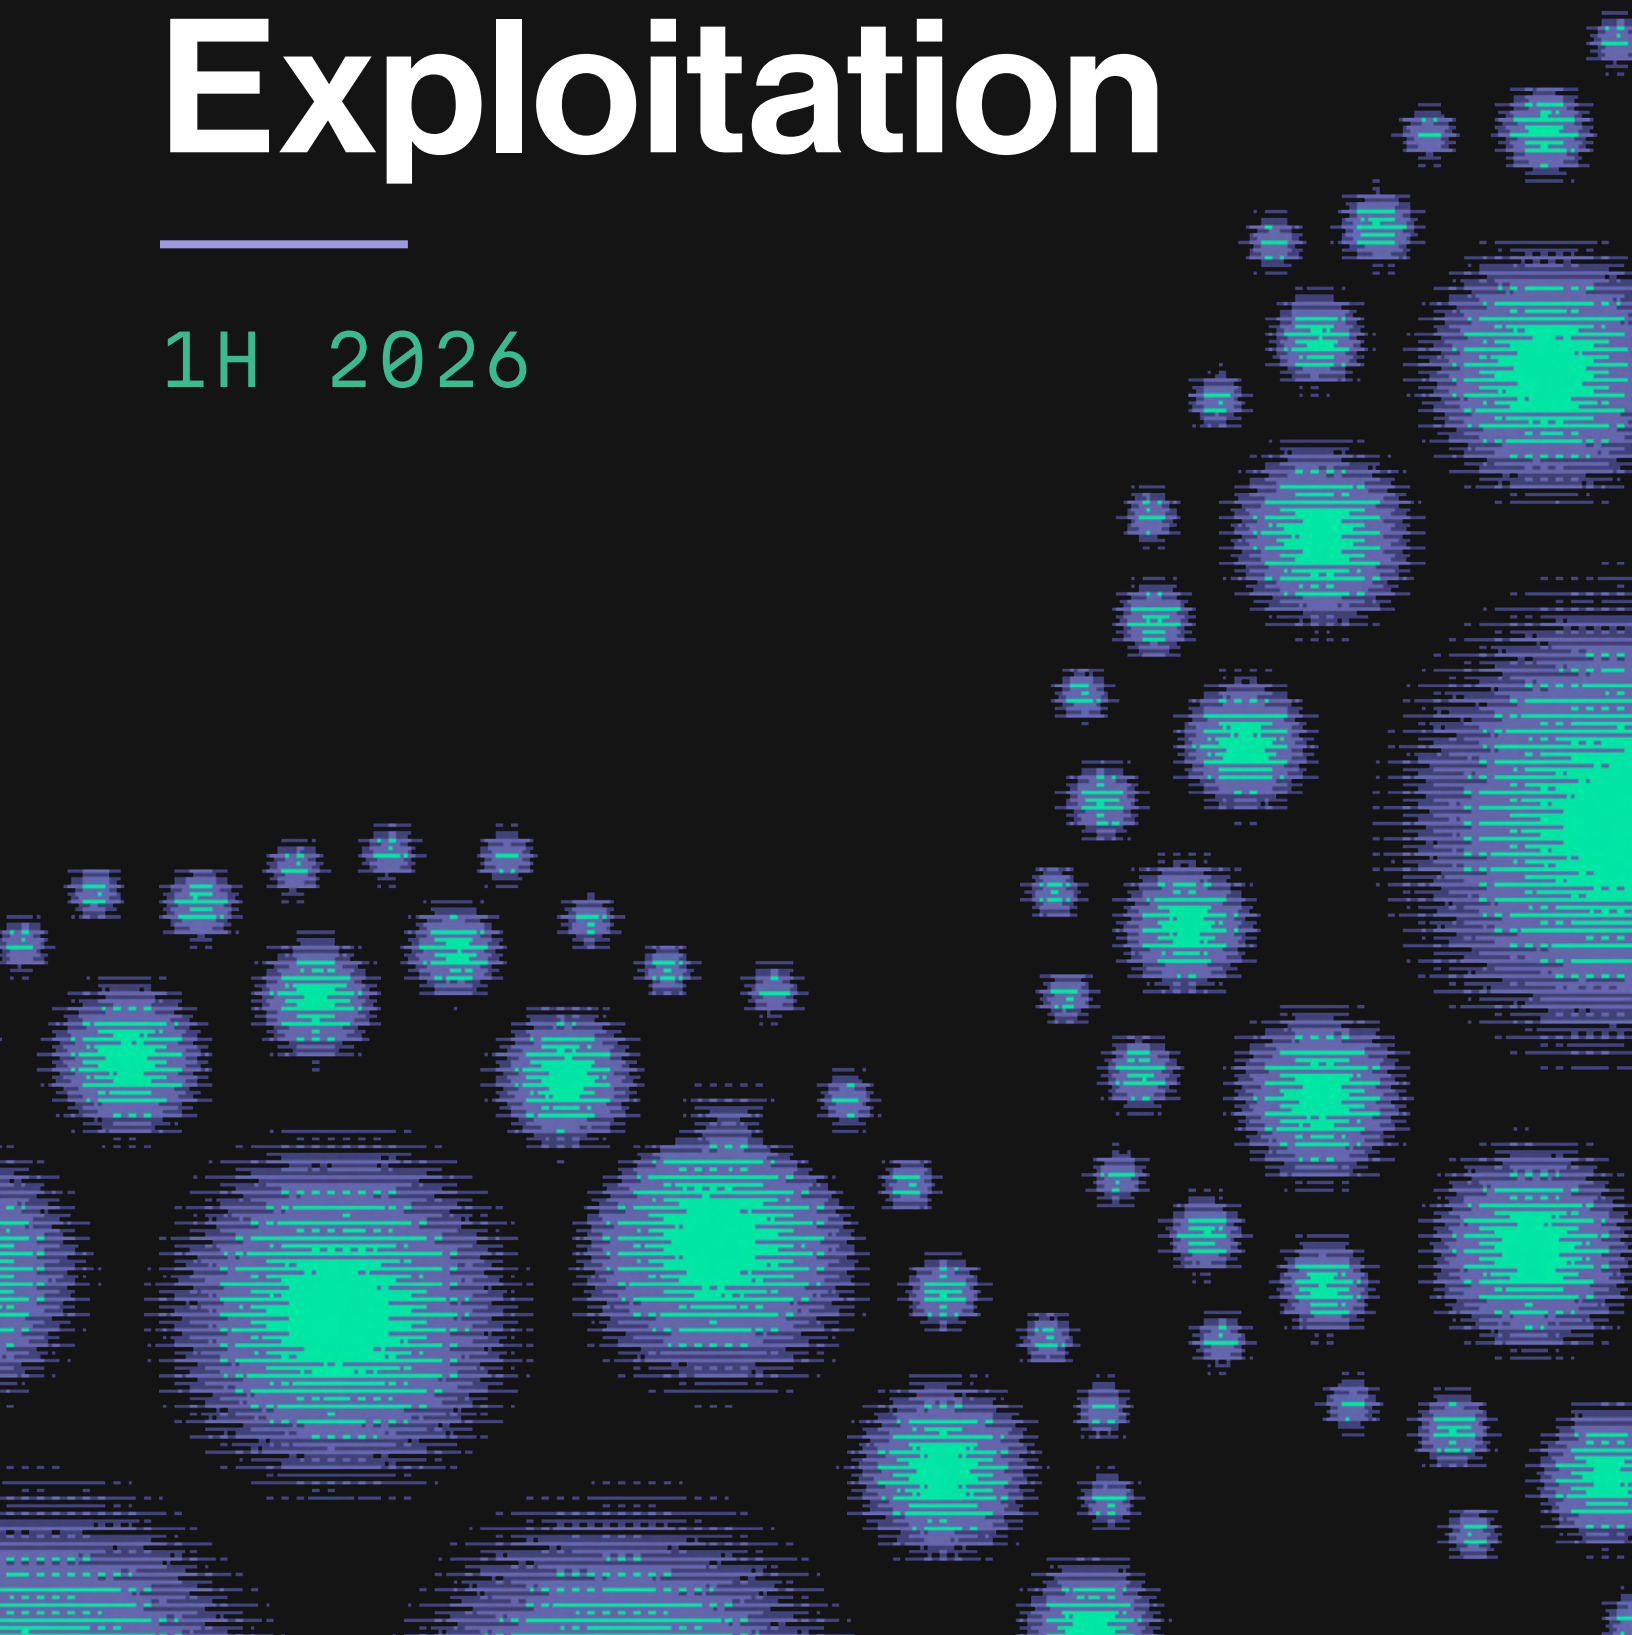

CONTENTS

Key Findings

Over the past six months, we've seen a **significant change in vulnerability discovery and disclosure** resulting in a substantial increase in the number of CVEs disclosed. This increase has come with warnings about the increase in vulnerability discovery by Autonomous AI systems, creating a "dangerous" scenario for the software ecosystem.

So, I was curious to explore:

- Are more vulnerabilities being discovered and disclosed, resulting in more vulnerabilities being exploited?
- Is the rate at which vulnerabilities are being exploited faster?
- Are vulnerabilities being discovered with AI tools more dangerous?
- Or are we all feeding into the AI-Assisted vulnerability discovery hype cycle?

01 EXPLOITED AT PUBLICATION

23.43%

In 2025, 28.93% of KEVs were disclosed as exploited on or before the day the CVE was published. That dropped to 23.43% in the first half of 2026.

02 TIME TO KNOWN-EXPLOITED

120→80 days

At the same time, vulnerabilities appear to be being exploited faster, with the median time from CVE publication to KEV falling from 120 days in 2025 to 80 days.

03 EARLY EXPLOITATION

31 days

Exploitation activity early in the CVE lifecycle remained steady, with roughly 200 CVEs becoming exploited within 31 days in the first half of 2026. Early exploitation activity is not signaling that it's scaling at the same pace as CVE issuance.

04 TARGETED TECHNOLOGIES

1/3 of KEVS

Content management systems remained the most targeted technology category, accounting for one-third of all KEVs - a more significant percentage of KEVs than we've seen historically.

05 AI-ASSISTED DISCOVERY

1,061

vulnerabilities attributed to AI-assisted discovery.

1.3%

confirmed exploited in the wild.

Of 1,061 vulnerabilities attributed to AI-assisted discovery, only 14, or 1.3%, have been confirmed as exploited in the wild, roughly matching the overall exploitation rate of all vulnerabilities in the first six months of the year.

---

06 NEW ATTACK SURFACE

AI GATEWAYS

WORKLOAD-SCALING PLATFORMS

MODEL-BUILDING TOOLS

WORKFLOW AUTOMATION

AI products are emerging as a new attack surface, with known exploitation affecting model-building tools, workload-scaling platforms, AI gateways, agents and workflow automation.

---

07 PROJECT GLASSWING

23k+

Reported findings

126

Published CVEs

1

Exploited in the wild

Anthropic reported more than 23,000 findings through Project Glasswing, but only 126 have resulted in published CVEs, and just one has been confirmed as exploited in the wild.

01

How fast are vulnerabilities being exploited in 2026?

TIME FROM CVE TO EXPLOITATION EVIDENCE

VulnCheck KEV 2026

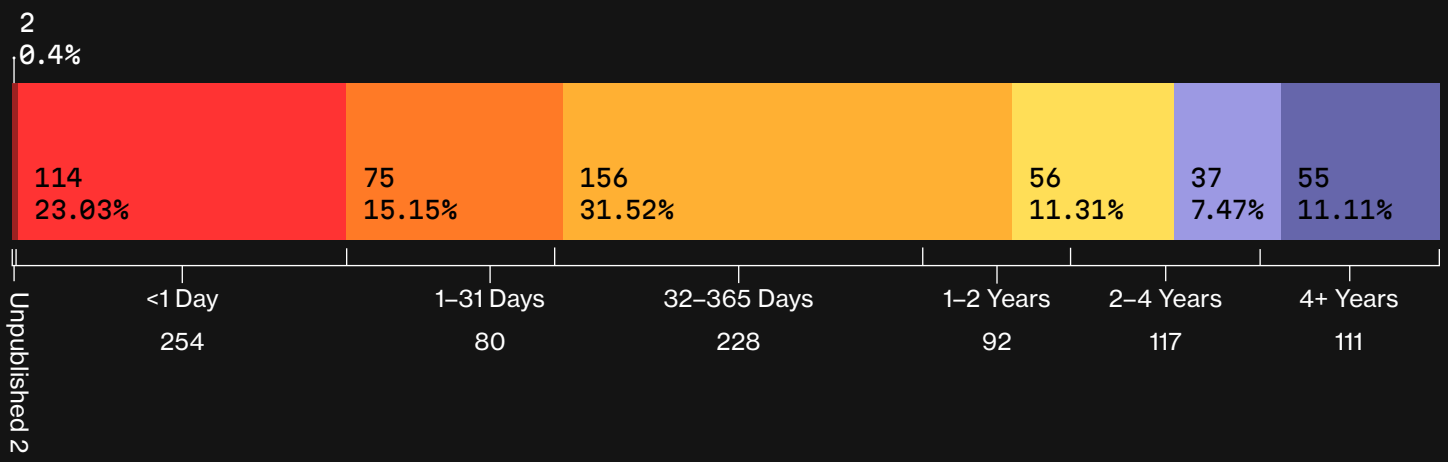

We started our research by looking at the rate at which vulnerabilities are being exploited. This analysis includes every Known Exploited Vulnerability added to **VulnCheck KEV** during the first half of 2026, based on CVE publication date and earliest evidence of exploitation. We identified 495 KEVs during that period.

In 2025, **28.93% of KEVs** were disclosed as exploited on or before the day the CVE was published. That figure dropped to 23.43% in the first half of 2026. At the same time, the median time from CVE publication to KEV fell from 120 days to 80 days, meaning that while there is slightly less evidence of exploitation on or before a CVE being published, vulnerabilities are reaching KEV status much faster overall.

With 23.43% of KEVs exploited on or before the day the CVE is issued, and a median of just 80 days to reach KEV status, it's clear why **CISA issued updated guidance in BOD 26-04**. The directive recommends prioritizing remediation based on risk and patching as aggressively as within three days when there's evidence of exploitation, automatability, high technical impact, and/or public exposure.

Is the rate of vulnerability exploitation increasing at the rate of disclosure?

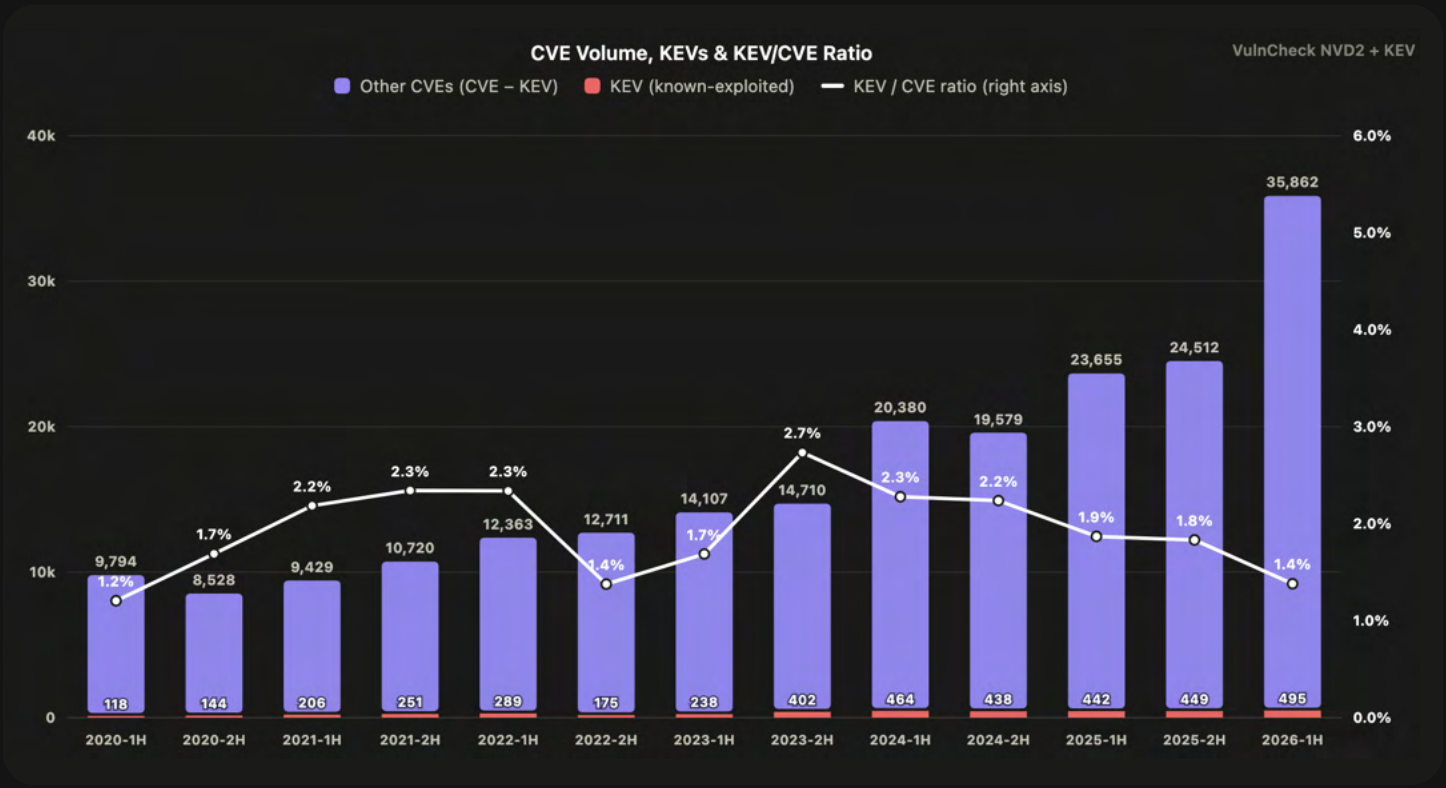

Looking at the past several years, vulnerability volumes have increased substantially, so we wanted to explore whether known exploited vulnerabilities were growing at a similar rate. We looked at the data in several ways.

The chart above shows the ratio of new KEVs added to VulnCheck KEV compared to the number of CVEs published during the same period. As both CVE and KEV volumes have continued to increase, the KEV-to-CVE ratio has declined over the past two years, peaking at 2.7% in the second half of 2023 and falling to 1.4% in the first half of 2026.

While the first half of 2026 saw a 10% increase in KEVs compared to the prior six months, CVE volume grew at a much faster rate of 45%, resulting in a significant drop in the KEV-to-CVE ratio. Of course, exploitation often occurs months or even years after a vulnerability is disclosed, so it's still too early to determine whether exploitation volumes will eventually follow the same growth trend as CVE issuance or level off at current rates. We'll have to wait and see how publicly available Frontier AI Cyber models continue to progress over the next year.

Is exploitation evidence keeping pace with new CVEs?

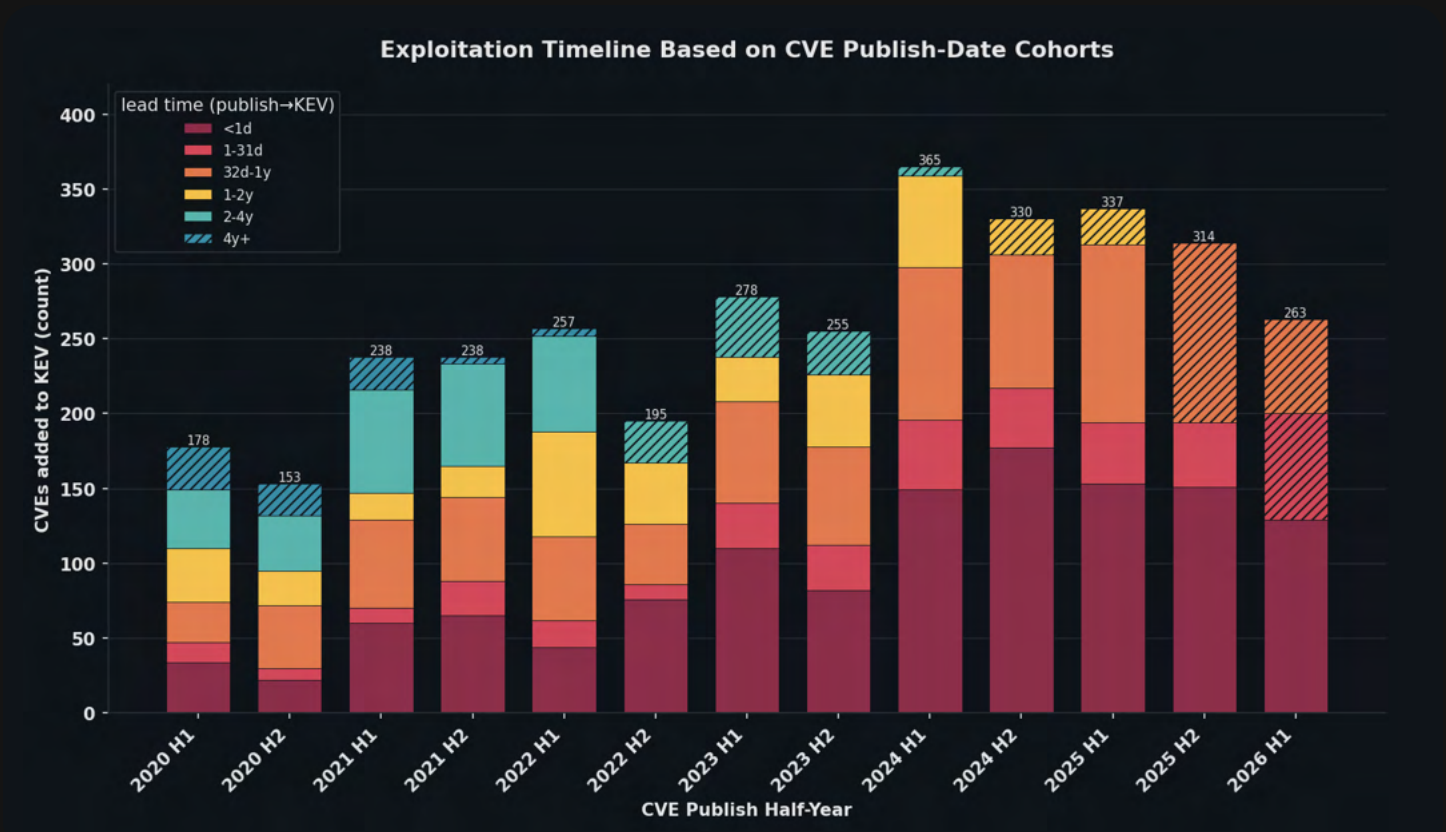

Next, I performed a cohort analysis looking at the publication date of the CVE for every KEV, grouped into half-year cohorts going back to 2020. This provides a view into long-term trends in vulnerability exploitation based on when CVEs were issued. The significant increase in exploited CVEs beginning in 2024 appears to closely align with the rapid growth in WordPress plugin-related KEVs after Patchstack, Wordfence, and WPScan began scaling CVE assignment for WordPress plugins.

The diagonally shaded segments represent lead time buckets that have not yet fully matured, so those bars will likely continue to grow as additional vulnerabilities are confirmed to have been exploited and new cohorts are added over time. Early exploitation activity in the first half of 2026 is holding steady rather than accelerating: roughly 200 CVEs reached KEV within 31 days of publication in 2026 H1, in line with 2024 (196) and 2025 (194). But because CVE issuance has grown so sharply over the same period, that same activity represents a smaller share of published CVEs, meaning early exploitation is not scaling at the same pace as CVE issuance. That said, this is still early analysis. Exploitation evidence often surfaces well after a CVE is published, so it will take time before we fully understand how recently published CVEs contribute to future exploitation trends.

04

What technologies are being targeted?

NUMBER OF KEVS BY TECHNOLOGY CATEGORY

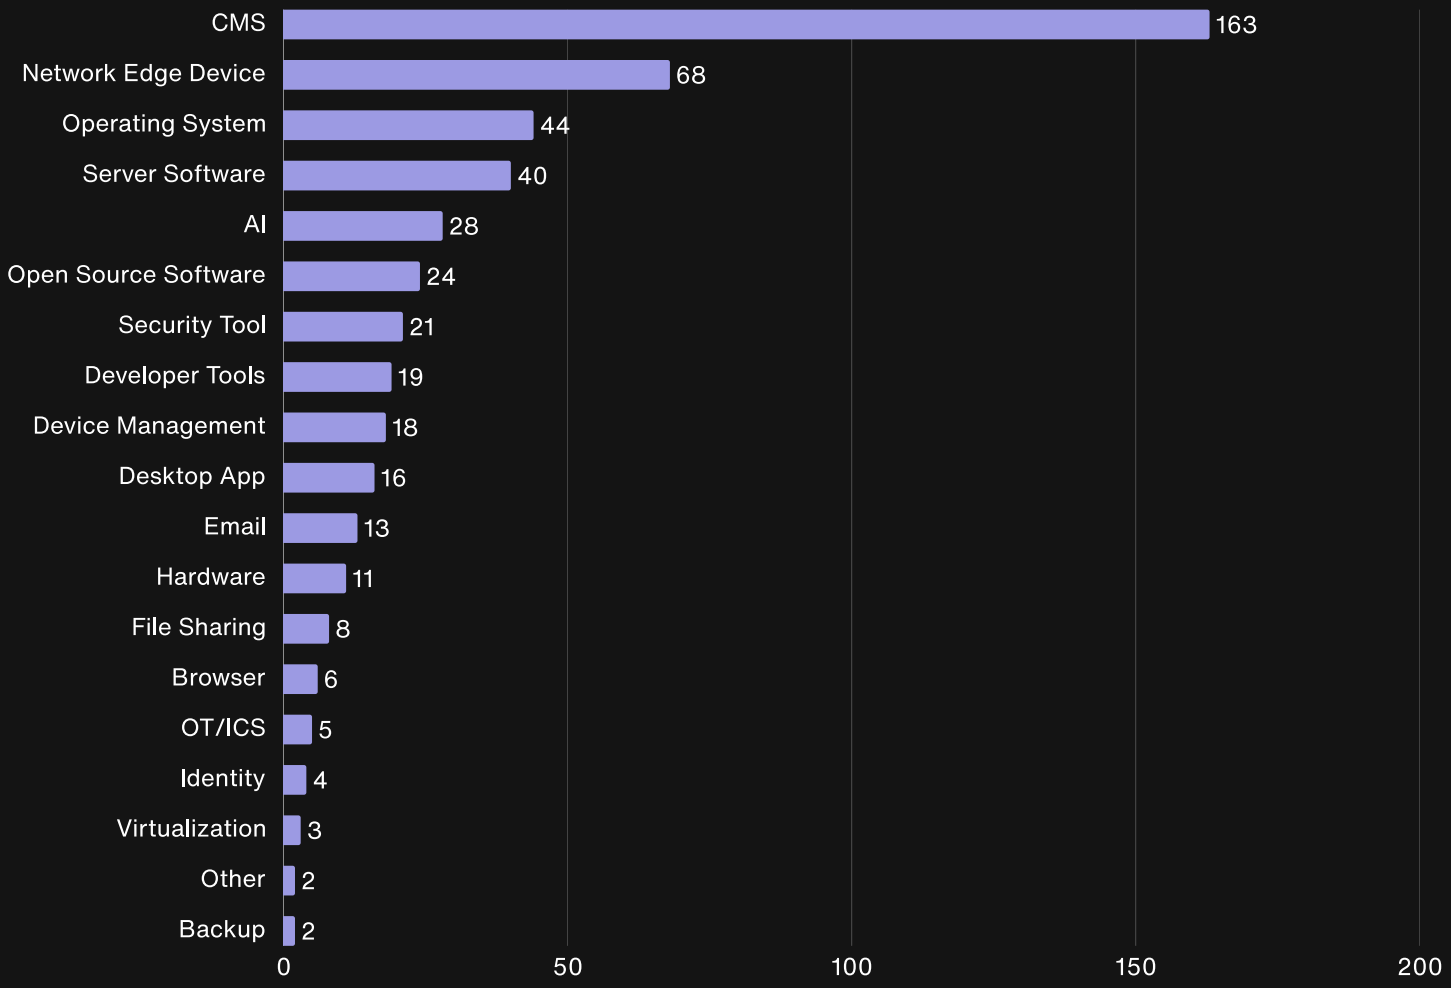

When we look at the top technology categories being targeted by exploitation activity, Content Management Systems (CMS) stand out, accounting for one-third of all KEVs we've added in the first half of 2026. It's typical for CMS to rank as a top category, but this is a more significant percentage of KEVs than we've seen historically.

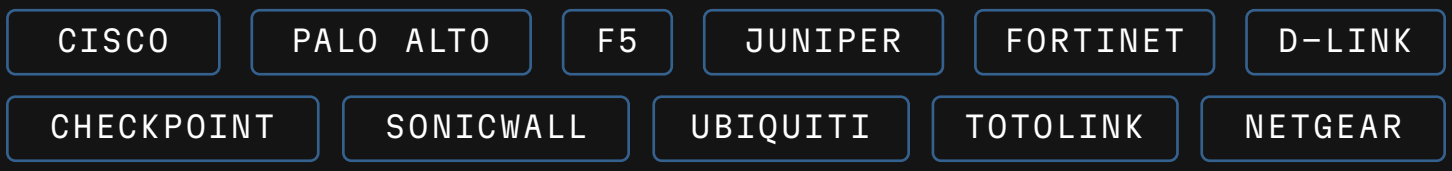

There has also been a recent advisory this month from the **Australian Signals Directorate on a large scale exploitation campaign targeting website content management systems**, which seems to align with the broader exploitation trend in CMS. While the large volume is associated with WordPress plugins, many CMS platforms have been targeted, including Drupal, Ghost, Kentico Xperience, and a handful of others. Given that CMS remains the single largest exploited category and shows no signs of slowing, organizations running any of these platforms should treat plugin and core patching as a standing priority.

Earlier this year, we researched Network Edge Devices in our **2026 State of Exploitation: Exploiting The Network Edge**. Network Edge Devices continue to be highly targeted, and the large majority of network edge device manufacturers continue to have new Known Exploited Vulnerabilities. Vendors that saw new KEV additions in the first half of the year include Cisco, Palo Alto, CheckPoint, F5, Juniper, Fortinet, SonicWall, Ubiquiti, TOTOLINK, Tenda, D-Link, Netgear, Linksys, and several others. It's also worth noting that during the first half of 2026, CISA issued **BOD 26-02, Mitigating Risk From End-of-Support Edge Devices**.

05

How quickly are the top 10 technology categories being exploited?

TOP TARGETED TECHNOLOGIES

Time from CVE to Exploitation Evidence

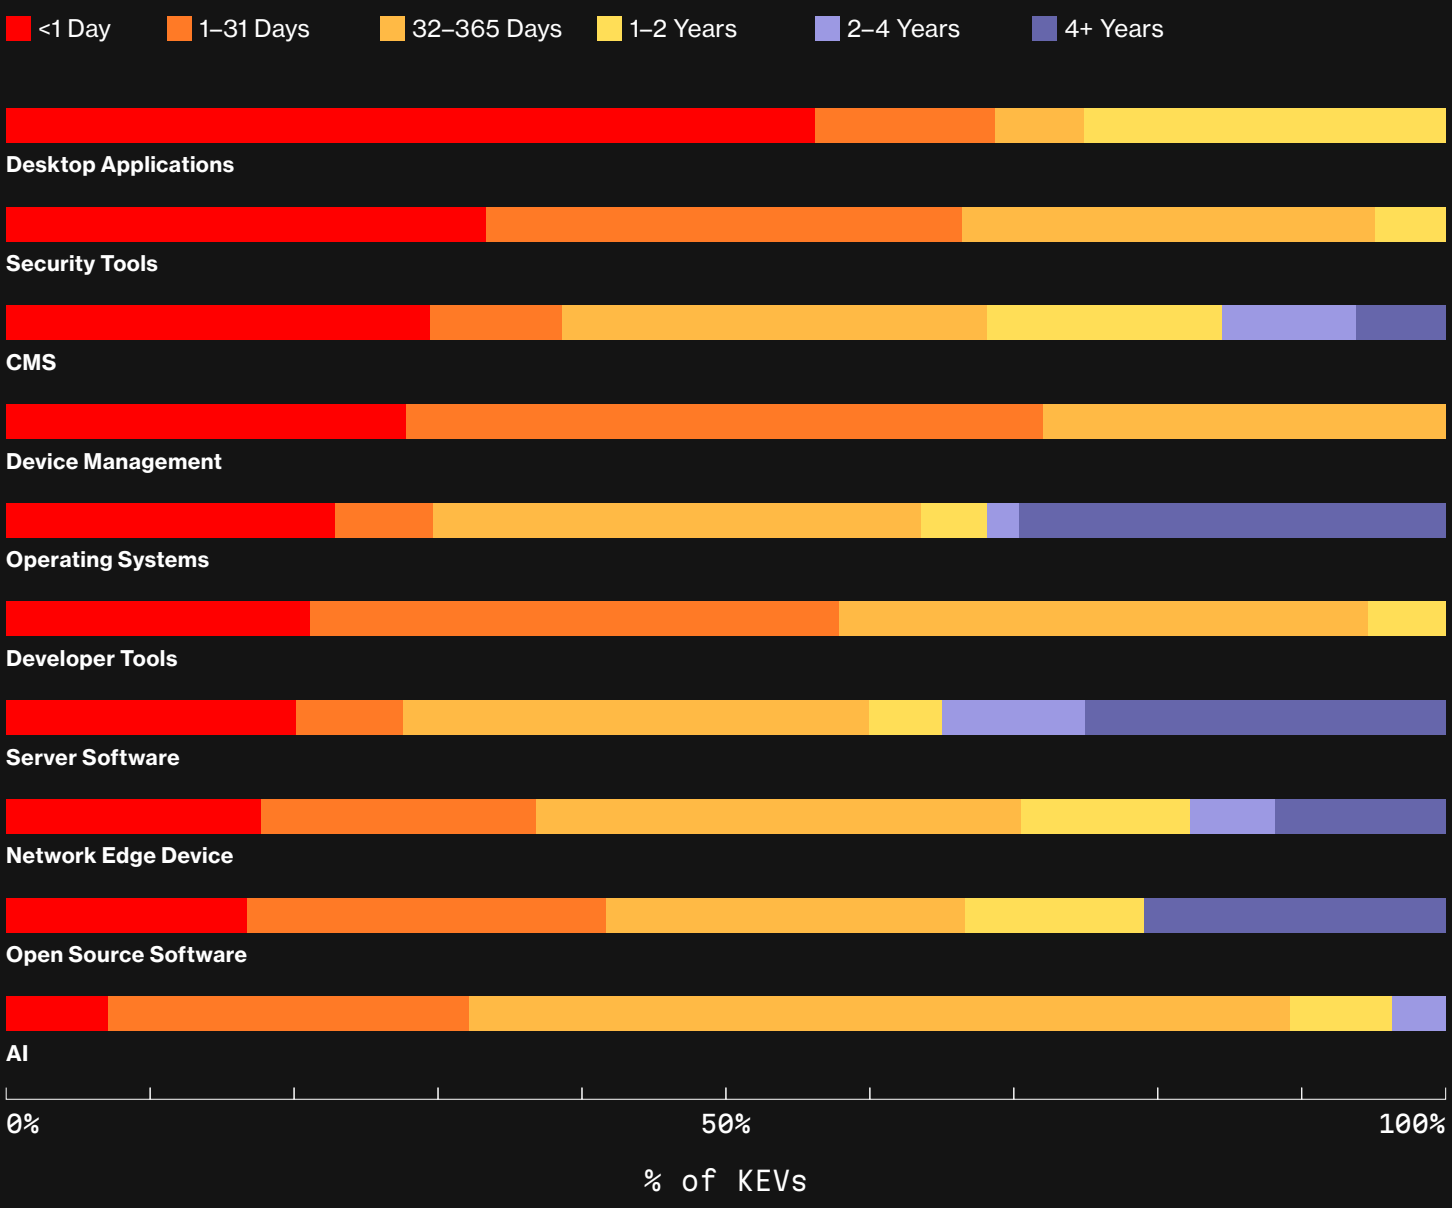

FFICE FORTINET EMS IVANTI EMM SOLARWINDS FORTSIEM SPLUNK

ER ADOBE ACROBAT NOTEPAD++ TREND MICRO APEX ONE WEAVER E-SUI

From a speed perspective, Security Tools (Splunk, Microsoft Defender, Trend Micro Apex One, Fortinet FortiSIEM, Fortinet FortiSandbox, Aqua Security Trivy, etc.), Developer Tools, Device Management Platforms (Ivanti EMM/Sentry, Fortinet EMS, SolarWinds, ConnectWise, Microsoft Configuration Manager, etc.), and Desktop Applications (Office, Adobe Acrobat, Weaver E-Suite, Notepad++, etc.) appeared to be exploited at a faster rate than other technology categories during the first half of 2026. This is likely due to the expansive deployments and access these technologies provide in Enterprise environments.

06

What AI technologies have been exploited?

AI PRODUCTS EXPLOITED

In 1H-2026

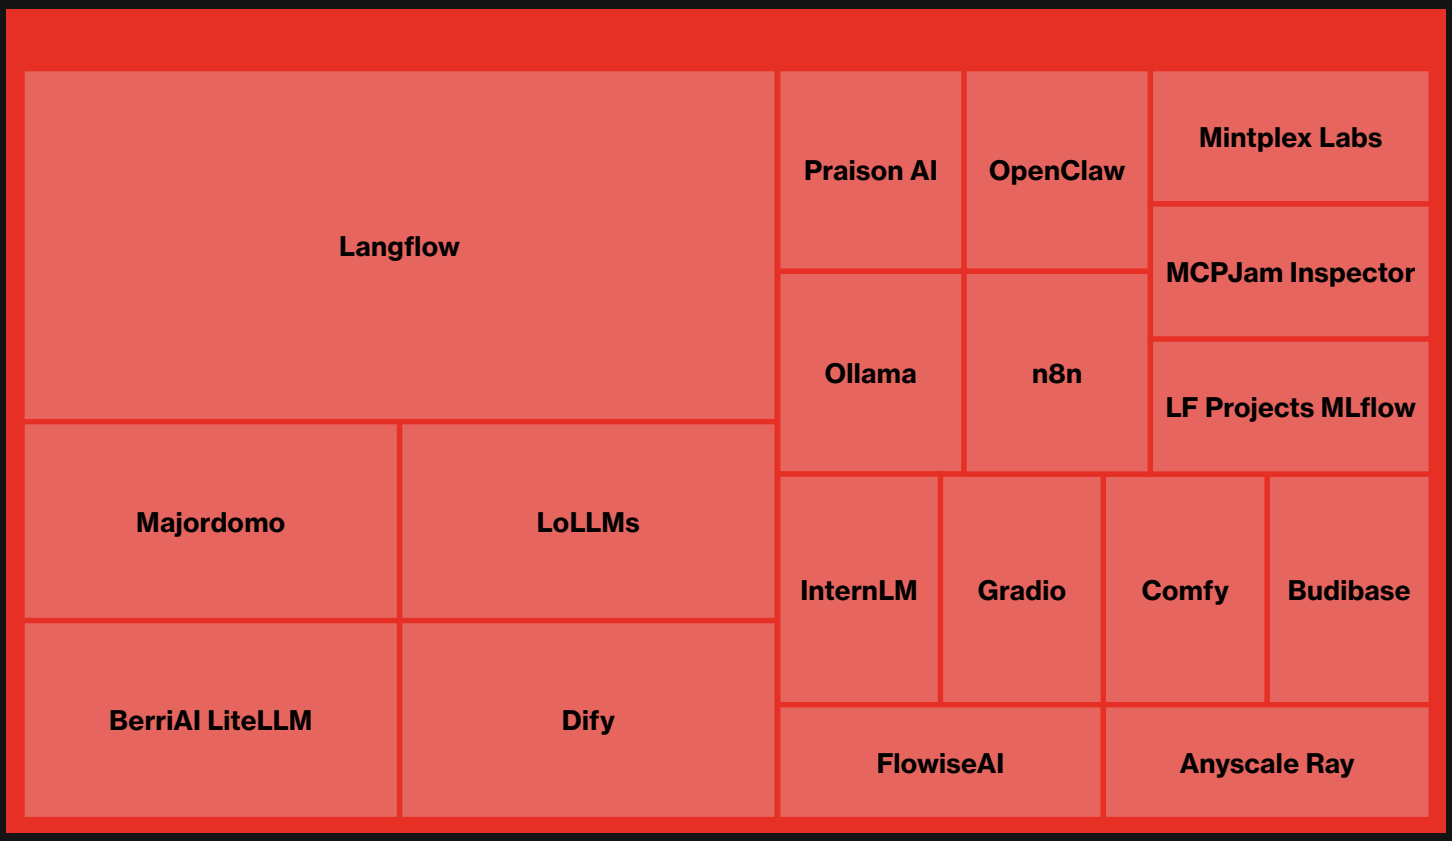

During the first half of 2026, we saw an increase in known exploited vulnerabilities targeting AI products. These products, many of which are open source projects, are used for model building, workload scaling, AI gateways, AI agents, AI workflow automation, and more. AI is increasingly being given access to sensitive corporate information as well as key infrastructure, often containing privileged information, and are often an easy target for Remote Code Execution (RCE), so it's logical that attackers would start experimenting with accessing these services.

We've observed exploitation activity across 10 of the 28 KEVs identified in AI systems in VulnCheck Canaries. With LangFlow, we've seen attackers gain initial access using exploits targeting both CVE-2026-0769 and CVE-2026-5027, harvest credentials, likely for services such as OpenAI and Claude, deploy cryptominers, and attempt lateral movement. Neither of these vulnerabilities have been added to CISA KEV.

07

What is the impact of AI-assisted vulnerability discovery?

PERCENTAGE OF EXPLOITED AI-DISCOVERED VULNERABILITIES

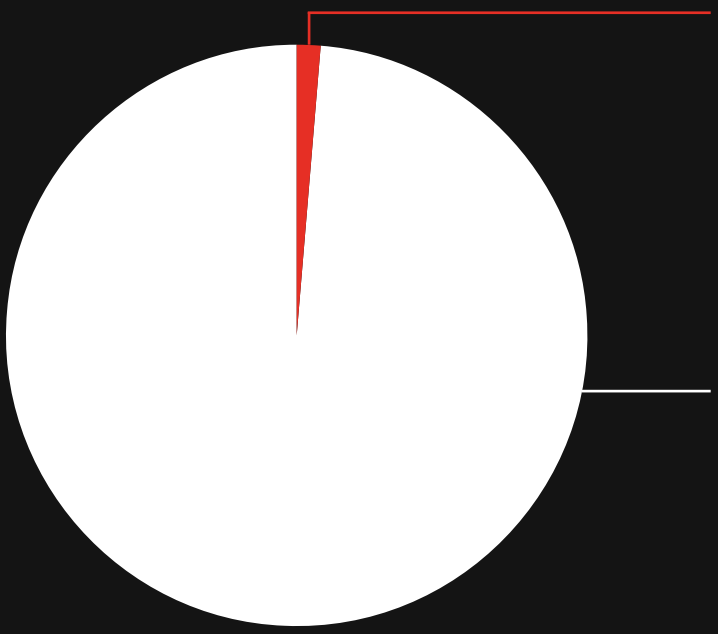

1.3%

of the 1,061 CVEs Identified by Berkeley Vulnerability Initiative and Patrick Garrity's Anthropic Credited CVEs project are known to be exploited.

98.7%

No Evidence of Exploitation

Recently, we've seen a growing emphasis on AI-assisted vulnerability discovery. In April, Anthropic announced Project Glasswing, warning that AI-assisted vulnerability discovery could enable attackers to hijack systems, disrupt operations, or steal data. I immediately began tracking disclosures attributed to Anthropic to understand whether any of the vulnerabilities they discovered were ultimately exploited in the wild.

While doing that, the **Berkeley Vulnerability Research Initiative** launched, providing another source for tracking AI-assisted vulnerability discovery. I consolidated both datasets and correlated them with VulnCheck KEV to better understand how often AI-assisted discoveries end up being exploited in the wild.

Across both datasets, 1,061 vulnerabilities are attributed to AI-assisted discovery. Of those, 14, or 1.3%, have been confirmed as exploited in the wild, and VulnCheck observed four of them being used against our canaries. That closely aligns with the exploitation rate we've observed across all vulnerabilities in the first half of 2026. While AI-assisted vulnerability discovery clearly has value for both attackers and defenders, the data does not suggest that AI discovered vulnerabilities are inherently more likely to be exploited than those found through traditional methods. Instead, AI appears to increase the volume of vulnerabilities that can be discovered, giving defenders an opportunity to identify and remediate them before attackers do. Time will tell whether the rate of exploitation increases as frontier cyber models become more widely available.

What products were exploited using AI-discovered vulnerabilities?

With the relatively small volume of AI discovered vulnerabilities with confirmed exploitation, our main observation is that it's early to come to any significant conclusions on what vendors / products are most likely to be exploited, however, it's worth noting that the KEVs include both commercial (Microsoft Windows, BeyondTrust) and open source products (ghost, chef, linux, etc.). It will be interesting to watch as additional evidence surfaces.

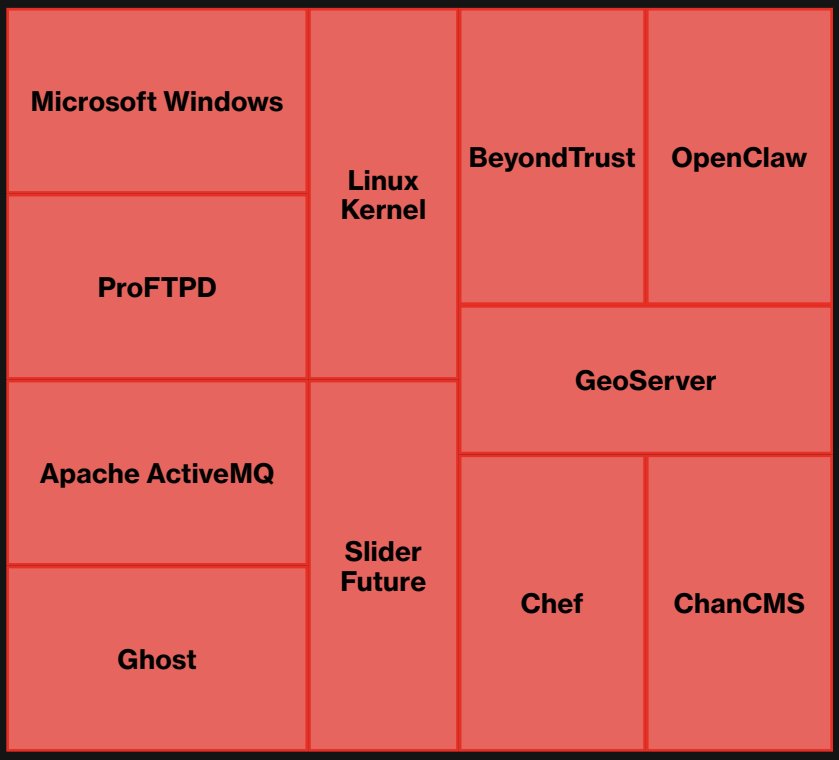

CVEs Identified by Berkeley Vulnerability Initiative and Patrick Garrity's Anthropic Credited CVEs project are known to be exploited.

09

Has Anthropic Glasswing lived up to the hype it brought this year?

ANTHROPIC MYTHOS LEDGER — CUMULATIVE FINDINGS OVER TIME

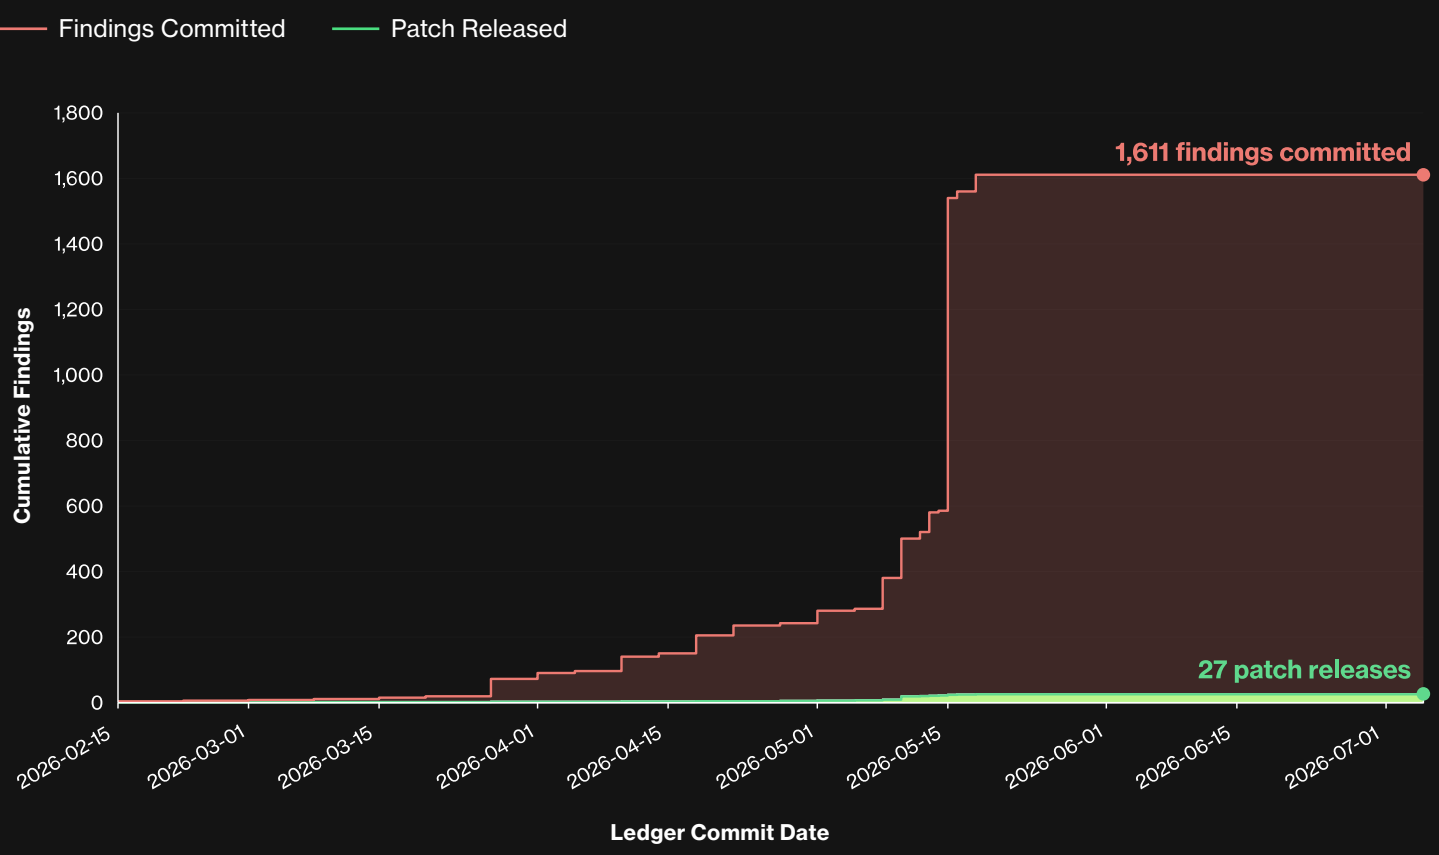

As mentioned earlier, Anthropic generated a lot of hype around the dangers of AI-assisted vulnerability discovery, so we also looked at that claim in isolation. If you recall, Anthropic launched a disclosure ledger in May, claiming Claude had identified 23,019 findings. What's surprising is that the disclosure ledger itself has never grown beyond the original 1,611 committed entries at launch, and Anthropic has failed to publish any updates or new disclosures despite more than 150 findings having passed their disclosure deadline as outlined in the Anthropic Coordinated Disclosure Policy. Based on the disclosure data I've collected, 126 of those findings have resulted in published CVEs, and just one of them, CVE-2026-26980, has been confirmed as exploited in the wild. At VulnCheck, we've observed first-hand exploits being used against this vulnerability on VulnCheck Canaries.

What KEVs in 1h-2026 are exposed on the internet?

INTERNET EXPOSURE BY UNIQUE KEV COUNT

Source: VulnCheck Target Intelligence / 1H-2026 VulnCheck KEV Additions

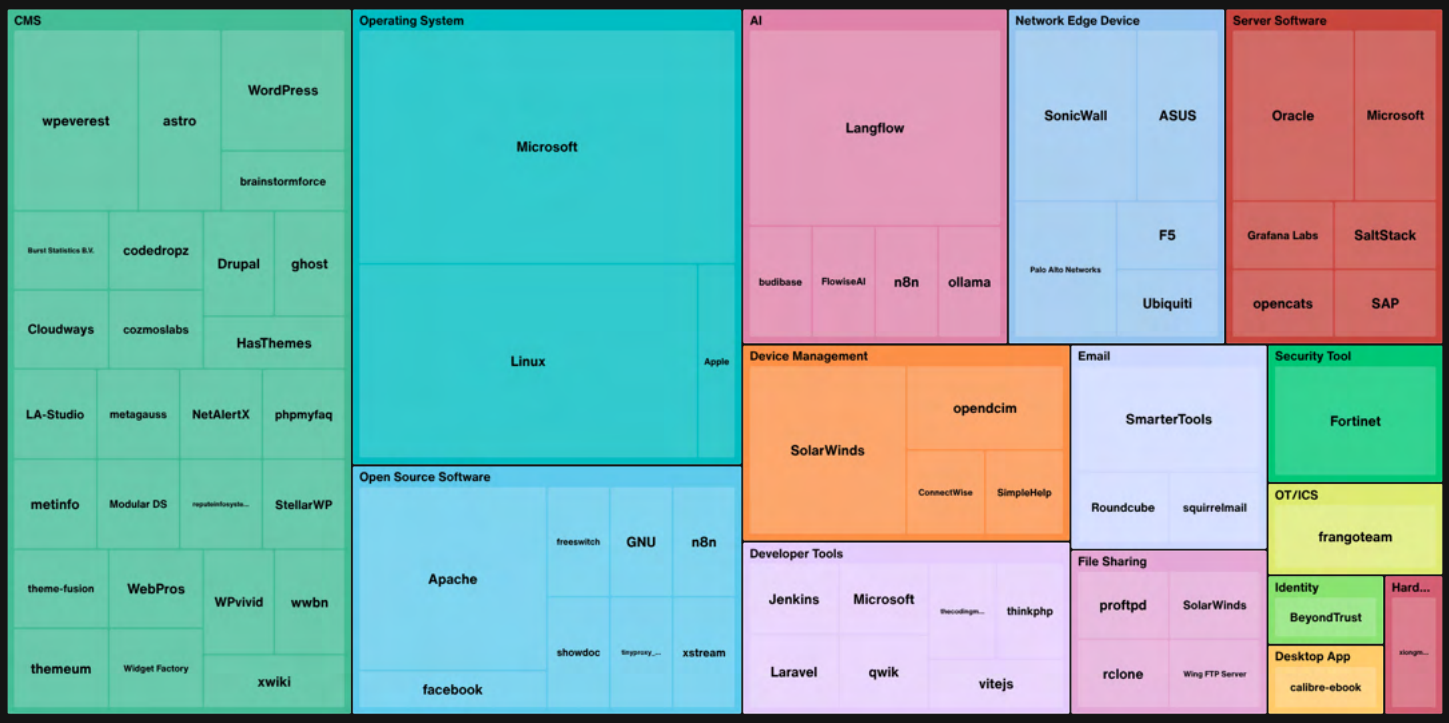

Expanding our view into exploited vulnerabilities that were first exploited in the first half of 2026, we explored VulnCheck Target Intelligence to better understand what vulnerabilities we were able to identify with hosts on the public internet that are still vulnerable. This provides visibility into the types of technologies that we know are both known to be discoverable on the internet and have vulnerabilities that have been exploited in the wild.

How did VulnCheck contribute to global visibility into exploitation?

495

CVEs were exploited in the wild in 1H 2026.

61

CVEs were observed by VulnCheck Canaries to have exploitation activity.

34

of the 495 KEVs identified as exploited were published by the VulnCheck CNA.

Only

12

have been added to CISA KEV.

During the first half of 2026, of the 495 CVEs that became exploited in the wild, we observed exploitation activity for 61 KEVs in VulnCheck Canaries, which are real vulnerable hosts. Only 12 of these CVEs have been added to CISA KEV.

During this time, we scaled CVE issuance. Our objective is straightforward: to ensure that vulnerabilities are discoverable by defenders and that exploitation evidence is available as early as possible to mitigate the risk of exploitation. 34 of the 495 KEVs identified as exploited were published by the VulnCheck CNA.

That's a primary reason we've invested in contributing back to the CVE program as a CVE Numbering Authority, providing our free vulnerability reporting service, and offering VulnCheck KEV as a free community service while partnering with third-party organizations and security researchers to ensure timely CVE issuance and access to evidence of exploitation.

Our objective is straightforward:

Ensure vulnerabilities are discoverable by defenders and that exploitation evidence is available as early as possible to mitigate the risk of exploitation.

Who was the first to report exploitation evidence?

EARLIEST REPORTER OF EXPLOITATION IN THE WILD

Source: VulnCheck KEV (1H-2026)

EARLIEST REPORTER OF EXPLOITATION IN THE WILD

Source: VulnCheck KEV (1H-2026)

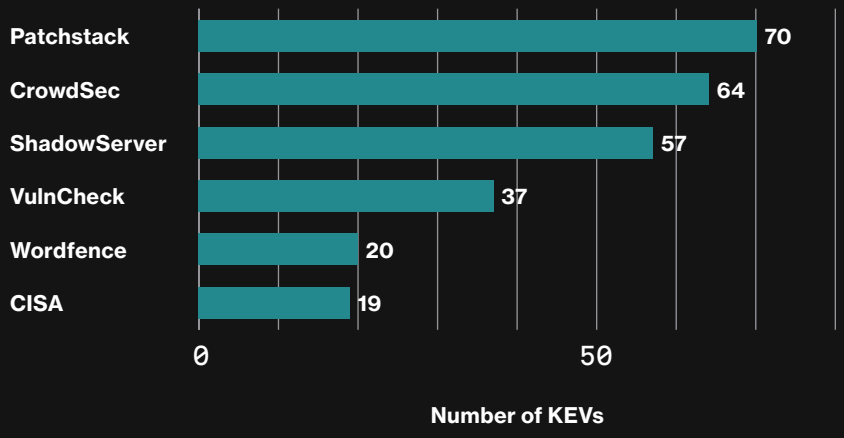

During the first half of 2026, 79 unique sources were the first to report exploitation. The top six sources first to report include Patchstack (70 KEVs), CrowdSec (64 KEVs), ShadowServer (57 KEVs), VulnCheck (37 KEVs), Wordfence (20 KEVs), and CISA (19 KEVs).

13

What to expect in an AI-assisted vulnerability era

In summary, the exploitation evidence we observed and collected in VulnCheck KEV signals that exploitation activity is sustaining from a volume and velocity perspective, and threat actors continue to target public-facing technologies such as CMS and Network Edge Devices.

At the same time, it appears that AI products and AI-assisted vulnerability discovery are contributing to a broader attack surface, and it will be interesting to observe the implications over time. For now, giving defenders access to more advanced frontier models is more likely to give defenders an advantage in strengthening software than to give attackers an advantage in discovering vulnerabilities before the software producers do.

The data so far, including Anthropic's own stalled disclosure ledger, suggests that AI-assisted vulnerability discovery and frontier capabilities have been overhyped relative to the evidence available today. That doesn't mean the risk is imaginary. It means the impact has been real but modest, and it's worth watching closely as the technology matures rather than reacting to the hype alone.

Patrick Garrity is a security researcher at VulnCheck where he focuses on vulnerabilities, vulnerability exploitation and threat actors.

Join the VulnCheck Community

Explore the leading resource for open vulnerability and exploit intelligence:

- VulnCheck KEV
- NVD++
- VulnCheck Exploit Database (XDB)

Join the Community →

VulnCheck

THE EXPLOIT INTELLIGENCE COMPANY