GLOBAL VULNERABILITY INTELLIGENCE REPORT 2026

DATA.	INTELLIGENCE.	ACTION.
48,000+	2215	256
VULNERABILITIES	EXPLOITABLE	EXPLOITED-IN-WILD

Executive Summary

2025 was a watershed year for cybersecurity. The convergence of AI-powered attacks, industrialized ransomware operations, and a shrinking vulnerability exploitation window created a threat environment that outpaced most organizations' defensive capabilities.

The Numbers

48,000+ CVEs were published, 131 per day, a 139% increase over five years. Of these, only **2,215 were confirmed exploitable**, and just 256 were exploited in real attacks. That means **99.5% of published CVEs were never exploited.** The signal-to-noise ratio collapsed from 1-in-79 in 2021 to 1-in-188 in 2025. Organizations patching by CVSS score are scaling effort to be in the numbers game and reduce risk on paper than reducing real risk.

The Window Is Gone

The time-to-exploit has been shrinking. According to the Zero Day Clock project[1], median time from vulnerability disclosure to first exploit fell exponentially from under 7 days in 2023 to just **4 hours in 2024.** By 2025, the exploitation time went **negative** implying the majority of exploited vulnerabilities were weaponized even before the disclosure. The average enterprise still takes more than a month to patch - creating a gap where attackers have a large window of opportunity and leverage, and defenders are structurally behind.

Who's Exploiting What

61 CVEs (23.8%) of exploited vulnerabilities in real attacks were attributed to named threat actors. **China-nexus groups dominated with 31 CVEs across 25 distinct clusters,** more than Russia, Iran, and North Korea combined. **33 CVEs were linked to ransomware operations**, with Clop leading at 8 zero-days. Severity proved an unreliable guide: **32 exploited CVEs were rated Medium or Low,** flaws that CVSS-based triage would have deprioritized entirely.

PUBLISHED	EXPLOITABLE	EXPLOITED	ZERO-DAYS
48K +	2,215	256	104
131 per day · +139% in 5 years	4.6% of published CVEs	0.53% → what actually mattered	40.6% → no patch when exploited

Dual nature of AI

AI is accelerating both offense and defense in measurable ways. Cyber attackers with AI can now reverse-engineer a vendor patch, identify the underlying vulnerability, and generate a working exploit in minutes - turning every security fix into an exploit. This capability is compressing the time-to-exploit curve toward zero and industrializing what was once a manual and specialist craft, lowering down the entry barriers for the attackers.

Defense is advancing too. In 2025, Google's Big Sleep discovered CVE-2025-6965[2], a critical SQLite zero-day that was known only to threat actors and about to be exploited. It marks the first documented case of an AI agent directly foiling a real-world zero-day attack. Similarly, in April 2026, Anthropic announced Claude Mythos Preview[3]- a general purpose language model which is strikingly capable at identifying and exploiting zero-day vulnerabilities.

The point to be noted is that Mythos preview was not trained to have these vulnerability discovery and exploit development capabilities Rather, they emerged as a downstream consequence of general improvements in code, reasoning, and autonomy. **Mythos exposes a deeper truth: the same model that finds a vulnerability can build the exploit for it.** Anthropic's own assessment states explicitly that the improvements making the model more effective at patching vulnerabilities are the same improvements making it more effective at exploiting them. This speaks volume about the dual nature of AI even when it's not purpose built for security. The frontier model, Mythos is withheld from general release and is instead deployed through Project Glasswing, a restricted initiative involving 50+ technology, cybersecurity, and infrastructure companies.

Defensive AI faces a fundamentally harder problem: it must uncover and test every possible vulnerability and exploit path while the attacker only needs one. The feedback is delayed (did we get breached months later?), noisy (real alert or false positive?), and inherently incomplete. **AI amplifies both sides, but it amplifies the side with cleaner feedback loops faster.** As Mythos class tools become public and reach wider deployment **in 2026, they will multiply the volume of CVEs that defenders must reduce the risk instantly and instant patching won't be the solution as patches break systems.** Enterprises will have to resort to mitigation controls such as IOCs/IOAs, policy/ configuration level changes, strengthening security controls to reduce the risk associated.

The Bottom Line. The huge gap between reducing average time-to-exploit and average time-to-patch is **not a metric to optimize, it is a structural failure**. Threat-informed prioritization focusing on the **~250 CVEs per year that actually get exploited** is no longer the best practice, it is the only way forward to proactively stop attacks.

MEDIAN TIME FROM DISCLOSURE TO FIRST EXPLOIT

Source: Zero Day Clock project

From a week's window to no window at all.

Most enterprise patch cycles are governed by change management processes designed for stability, not speed. Those processes made sense when exploitation took months. Now, They're an organizational liability when it takes hours. Patch management programs need to renegotiate their change approval workflows to avoid this structural constraint.

Shannon Lietz
Co-Founder & CEO, ThirdScore

IN THIS REPORT

Table Of Contents

Ten chapters tracing the 2025 vulnerability lifecycle from the 48,000+ CVEs published, down to the ~250 that mattered, and the playbook for the year ahead.

CHAPTER 01

Year-Over-Year Vulnerability Exposure - The Funnel Of Relevance

CVE volume exploded 139% in five years, from 20K in 2021 to 48K+ in 2025, but the number exploited in real world attacks barely moved. That's **~250 CVEs per year, every year**, regardless of how fast the top of the funnel grows. The exploitable pool hovered between 1,983 and 2,660 across the five-year window, a modest 12% shift. The machinery that converts a published CVE into a working exploit has a ceiling, and the industry hit it years ago.

The Relevance Graph

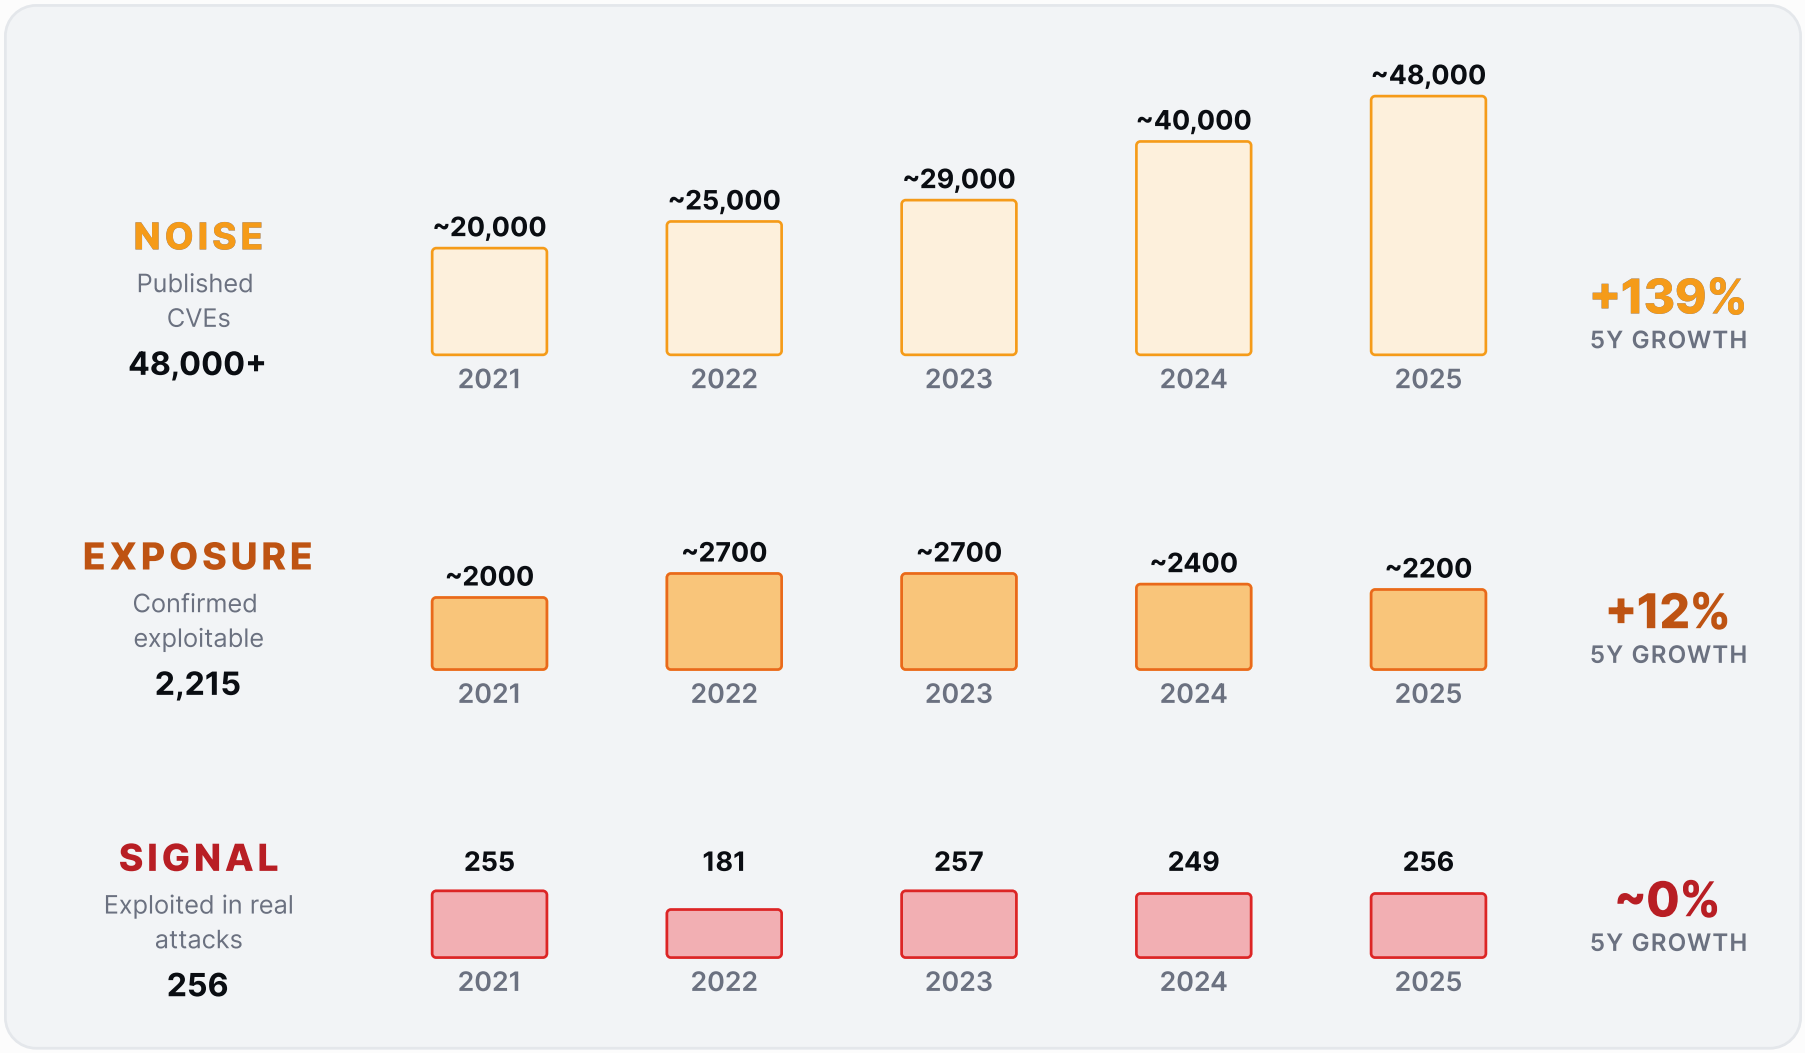

The Operational Cost of Chasing the Denominator

The signal-to-noise ratio didn't just worsen: it collapsed. In 2021, 1 in 79 published CVEs saw real exploitation. In 2025, it was 1 in 188. The funnel gets taller every year, but the bottom stays flat. We are drowning in **2.4× more noise** while the real threats stay hidden deeper.

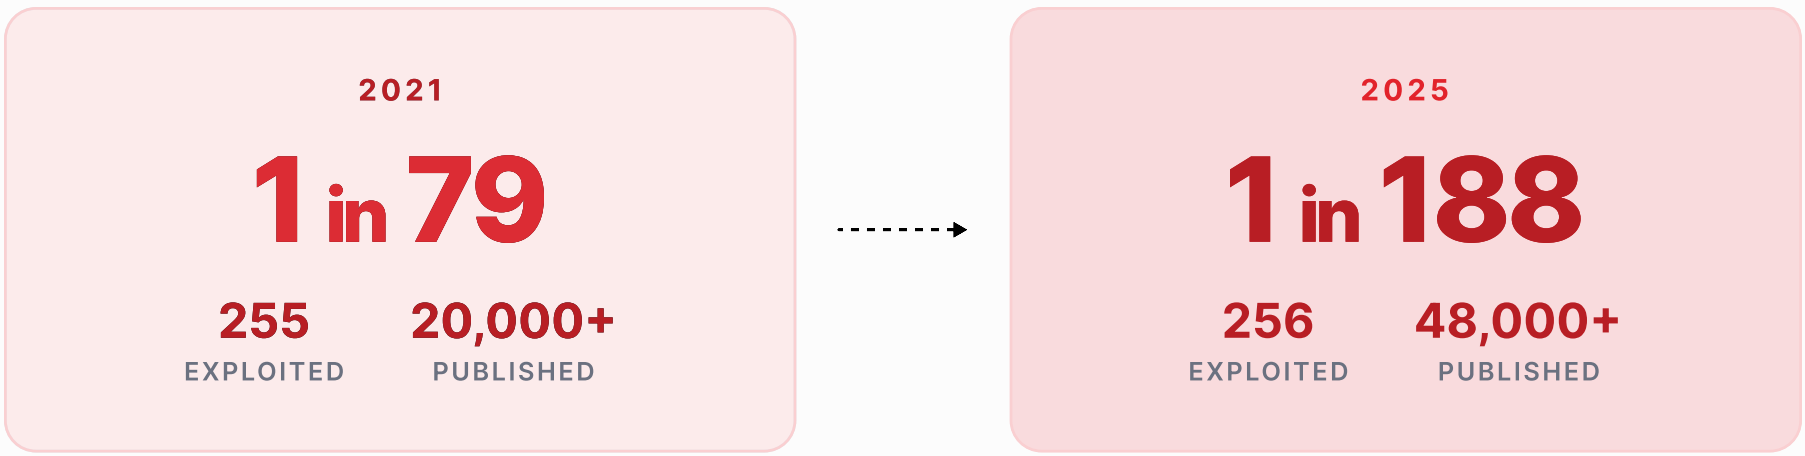

This has a direct operational cost. Organizations patching by volume or CVSS score alone are chasing a denominator that doubles while the numerator holds steady. For example, **in 2021 it was 255 / 20,000** and **in 2025 it's 256 / 48,000**.

Every increase in CVE output translates into more tickets, more scanning cycles, more patching windows spent on vulnerabilities no attackers are currently using. **99.5% of published CVEs in 2025 were never exploited in the wild.**

RECOMMENDATION

Focus on the vulnerabilities that actually matter. Out of the thousands of CVEs published each year, only around 250 are ever exploited in the real world. Without threat intelligence to spot those ones, teams end up patching by volume, chasing numbers instead of reducing real risk.

Use threat intelligence to prioritize the few vulnerabilities attackers are actively using, and patch those first.

CHAPTER 02

Exploitation Landscape

Of the **~48,000 CVEs published in 2025** - roughly 131 per day - only **6 per day had confirmed exploitability**, and fewer than 1 (0.7) was exploited in the wild. That final filter is ruthless: **99.5% of published vulnerabilities never saw real-world exploitation** and **95.5% of the CVEs aren't even exploitable**.

EXPLOITED IN THE WILD	ZERO-DAYS
256	**104**
0.53% of all CVEs	40.6% of exploited

RANSOMWARE	NAMED ACTOR
33	**61**
12.9%: mostly n-day	23.8% attribution rate

The 256 That Crossed the Line

The 256 that crossed that line tell a sharper story. **104 were exploited as zero days** before a patch even existed, giving defenders no remediation window. That's **40.6% of all exploited CVEs**, meaning for nearly half the threats that mattered, the traditional patch cycle was irrelevant from day one.

Of the 256, only 61 (23.8%) could be attributed to a named threat actor, and **33 (12.9%) were tied to ransomware operations**. The remaining majority, over half, were exploited without any known attribution: no known operator, no recovered malware.

40% of exploited CVEs in 2025 were zero-days. There was no patch. That makes the question of 'are we patched?' irrelevant. The more useful question organizations need to be able to answer is whether their existing controls would detect and contain the exploitation attempt with clear next steps of response and recovery - and most can't answer that with any confidence.

Rohit Parchuri
CISO at Yext

CHAPTER 03

Threat Actors & Malware

61 of 256 exploited CVEs (23.8%) had named-actor attribution. China-nexus groups dominated with 31 CVEs across 25 distinct clusters, more than Russia, Iran, and North Korea combined. **CVE-2025-8088 (WinRAR) became the universal weapon, used by 4 distinct nation-state groups.**

Attribution by Nation-State Nexus

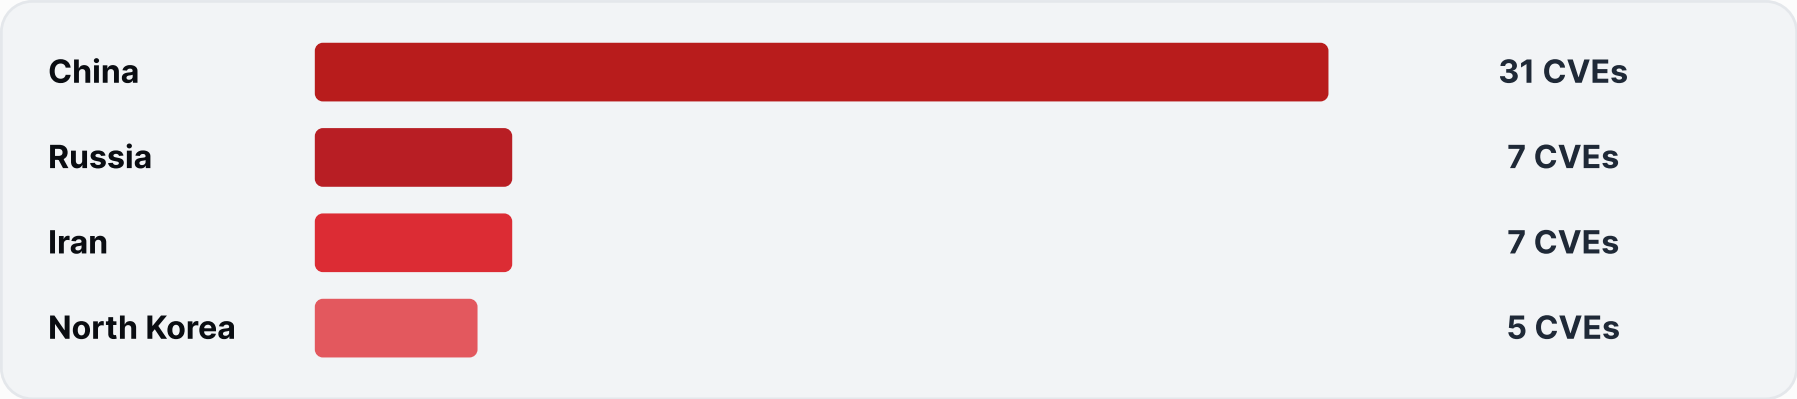

China isn't just the most active: it's operating at a fundamentally different scale. No other nation comes close. A single CVE can map to multiple groups across the same or different nexuses.

Sliver Replaces Cobalt Strike · Stealth Ops on the Rise

SLIVER · CROSS-ACTOR C2

8

8 CVEs · vs Cobalt Strike on 6

Sliver has replaced Cobalt Strike as the cross-actor C2 of choice. Sliver appeared on 8 CVEs (vs Cobalt Strike on 6), used by both China-nexus state actors AND ransomware groups. It's the new shared infrastructure layer - open-source, harder to signature, deployed by UTA0178, APT41, DragonForce, and Storm-1567 alike.

ATTRIBUTION · NO MALWARE

14

Living-off-the-land at scale

14 CVEs have actor attribution but zero malware recovered. Stealth operations from Volt Typhoon, UNC3886, Lazarus, Caramel Tsunami, operating without dropping detectable payloads. Living-off-the-land at scale.

CHAPTER 04

The Severity Paradox

High Severity ≠ High Risk; Low Severity ≠ Low Risk

Severity-based triage falls in both directions: it over-prioritizes and under-prioritizes real threats.

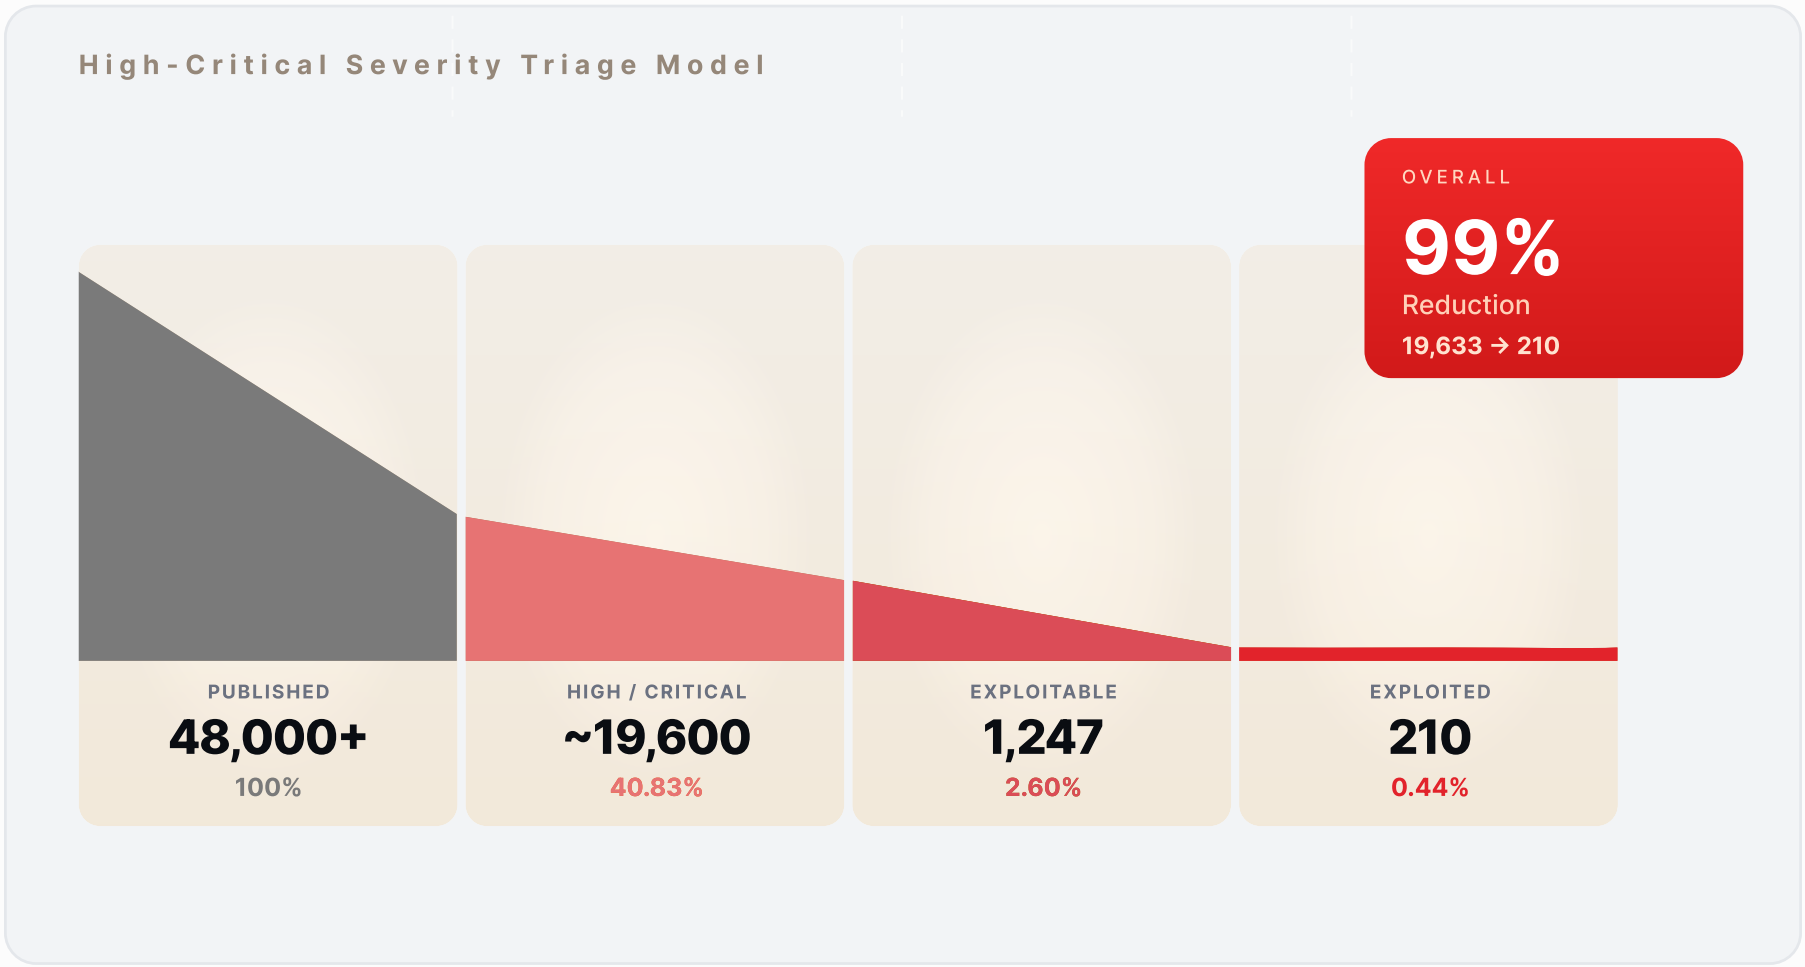

Over-Prioritization

It over-prioritizes, flagging **~19,600 Critical and High CVEs as urgent when only 210 were actually exploited**. Only exploitation evidence (not severity scores) separates the 210 that mattered from the 19,423 that didn't.

Under-Prioritization

It also **under-prioritizes 32 exploited CVEs that were rated Medium or Low**: flaws that CVSS-first triage would have placed at the bottom of the queue or skipped entirely. These weren't theoretical risks: they were actively weaponized by nation-state actors and ransomware operators while defenders focused elsewhere.

VOICE FROM THE FIELD

> ***"Attackers don't filter by severity; they filter by utility."***

Ankit Mani
Head of Hiveforce Labs

Recommendations - Multi-Step Prioritization

Vulnerability prioritization should be a multi-step method: exploitation evidence first, exploit availability second, severity and asset context after.

RECOMMENDATION #1

Vulnerability Prioritization Should Be a Multi-Step Method

Actively Exploited

The vulnerabilities that are currently getting exploited in the wild in active threat and malware campaigns and have been exploited in the past. These vulnerabilities represent imminent threat to your organization.

Exploitable but Not Yet Used

Identify the ones which can be exploited: for which exploits are already available but haven't been used in active campaigns yet.

Severity and Impact

Then prioritize based on the severity and impact of the underlying flaw. By this point, the actively dangerous set has been handled: severity is a refinement, not the primary signal.

RECOMMENDATION #2

Asset Criticality & Business Risk

When dealing with an overwhelming set of prioritized vulnerabilities, take asset criticality into account to calculate the business risk and prioritize based on the business risk and impact.

CHAPTER 05

Zero days : Root Causes & Patch Release

Zero-day vulnerabilities represent the most critical class of security threat: exploited before vendors can issue patches, leaving organizations with no immediate defense. Of the **256 CVEs exploited in the wild, 104 (40.6%) were zero-days**: meaning for nearly half the threats that mattered, the traditional patch cycle was irrelevant from day one.

ZERO-DAYS
104
40.6% of exploited CVEs

ATTRIBUTED
32%
Linked to known threat actor

TOP VENDOR
Microsoft
23 zero-days · top vendor

MALWARE TIED
47%
Identifiable malware deployment

TOP AFFECTED VENDORS · ZERO-DAY COUNT · 2025

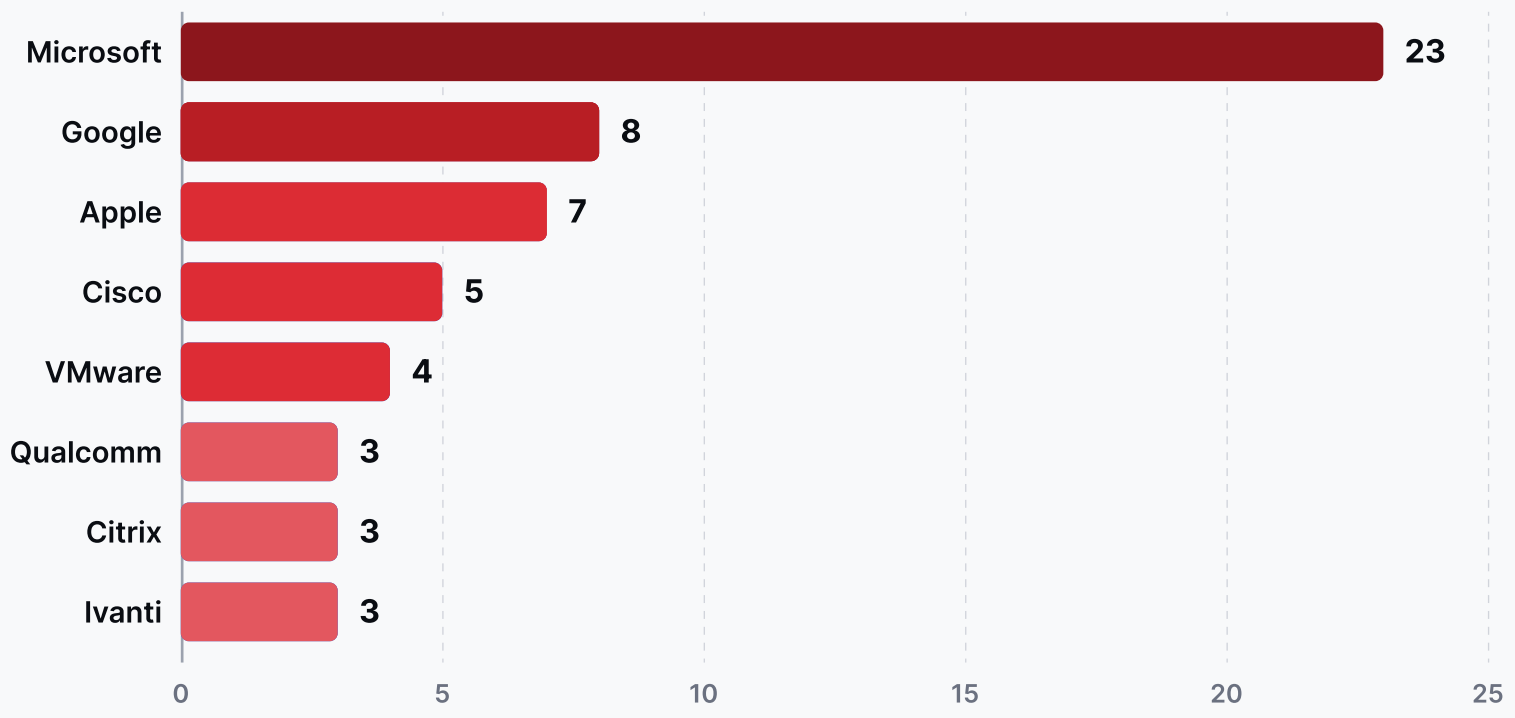

Chrome had eight zero-days in 2025. Browser-based exploitation chains are often treated as a user education problem; don't click bad links. They're an infrastructure problem as well. The browser is the most privileged application most users run, it's updated independently of enterprise patch cycles, and many exploitation requires nothing more than loading a page.

Poornaprajna Udupi
Co-Founder, CTO at Vinyl Equity

Root Cause of Zero Days

Understanding the root causes of zero-day vulnerabilities is essential for building effective defenses.

- **CWE-416 (Use After Free) is the most exploited weakness class with 12 instances**, followed by
- **CWE-20 (Improper Input Validation) with 7 and CWE-787 (Out-of-bounds Write) with 6.**
- **Notably, 17 of the 25 CWEs in MITRE's 2025 Top 25 list are represented in this dataset, while 31 CWEs fall outside the Top 25: indicating emerging attack vectors that traditional prioritization may miss.**

Zero Day Patch Release Time

The number of days between confirmed exploitation and vendor patch release - is a critical metric for understanding organizational risk exposure. Of the 104 zero-days, **45 received same-day or pre-disclosure patches** (largely from vendors who discovered exploitation internally), their first exploitation details largely remain unknown, leaving 59 CVEs with measurable response windows.

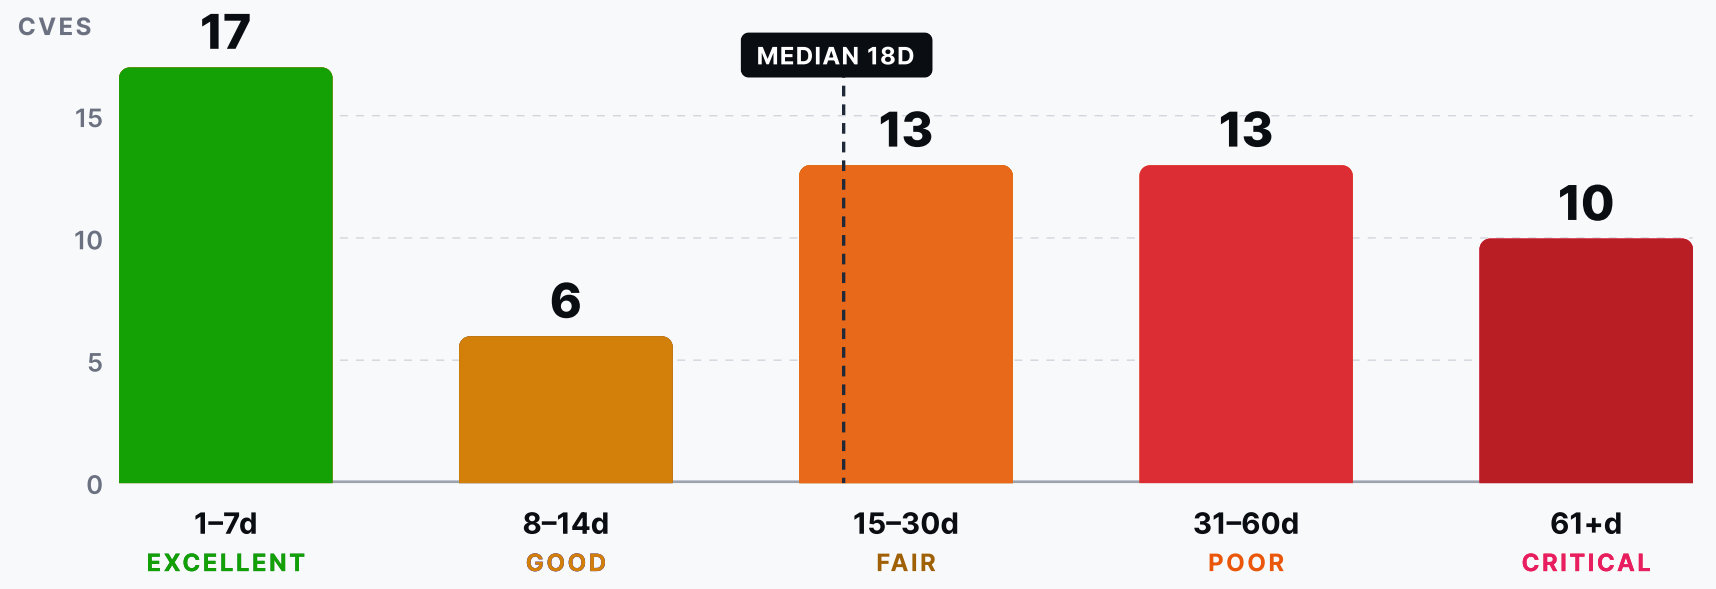

Excludes 45 CVEs patched same-day or pre-disclosure (window ≤ 0). Mean response 51.2 days; median 18 days; standard deviation 82.7 days.

61% of zero-days patch were released within 30 days; However, the **remaining 39% represent extended exposure window that attackers actively exploit**. Microsoft demonstrated the most consistent patch release cadence for its high-volume portfolio.

Several smaller vendors (particularly in IoT and enterprise appliance categories) exhibited release times exceeding 100 days. **The wide standard deviation of 82.7 days across all vendors indicates highly inconsistent patch release cadence industry wide.**

RECOMMENDATION #1

In the Absence of Patches

Reduce the risk of exploitation by strengthening existing security controls and putting compensatory controls in place. **Compensatory controls are your only defense against zero days**.

RECOMMENDATION #2

IoA Detection Beyond Signatures

Map real-world threat actor tradecraft and **create IoA detection logic derived from threat intelligence** to hunt malicious activities beyond the limitation of known signatures and available patches.

RECOMMENDATION #3

Validate Against Real Threats

Regularly perform Adversarial Exposure Validation (AEV) to check efficacy against real-world threats and measure detection and response capabilities.

CHAPTER 06

Where Exploited Vulnerabilities Sit in The Enterprise Network

The platform-level analysis of exploited-in-the-wild vulnerabilities paints a concerning picture of where real-world attacks are landing across enterprise infrastructure.

Security Platforms · 39 Flaws (15.2%)

Security platforms and solutions represent the single largest exposure, accounting for 39 flaws (15.2%): an ironic reality where the very tools deployed to defend the network become the entry point. This includes firewalls, VPN gateways, EDR/AV solutions, SIEM platforms, and IAM/PAM systems from vendors like Fortinet, Palo Alto Networks, Ivanti, SonicWall, and CyberArk.

Collaboration · 25 Flaws (9.77%)

Collaboration and communication platforms account for 9.77% (25 flaws): Zimbra, Microsoft SharePoint, RoundCube, and TeleMessage were actively targeted by nation-state actors seeking access to sensitive internal communications.

WHERE EXPLOITED CVE LIVES

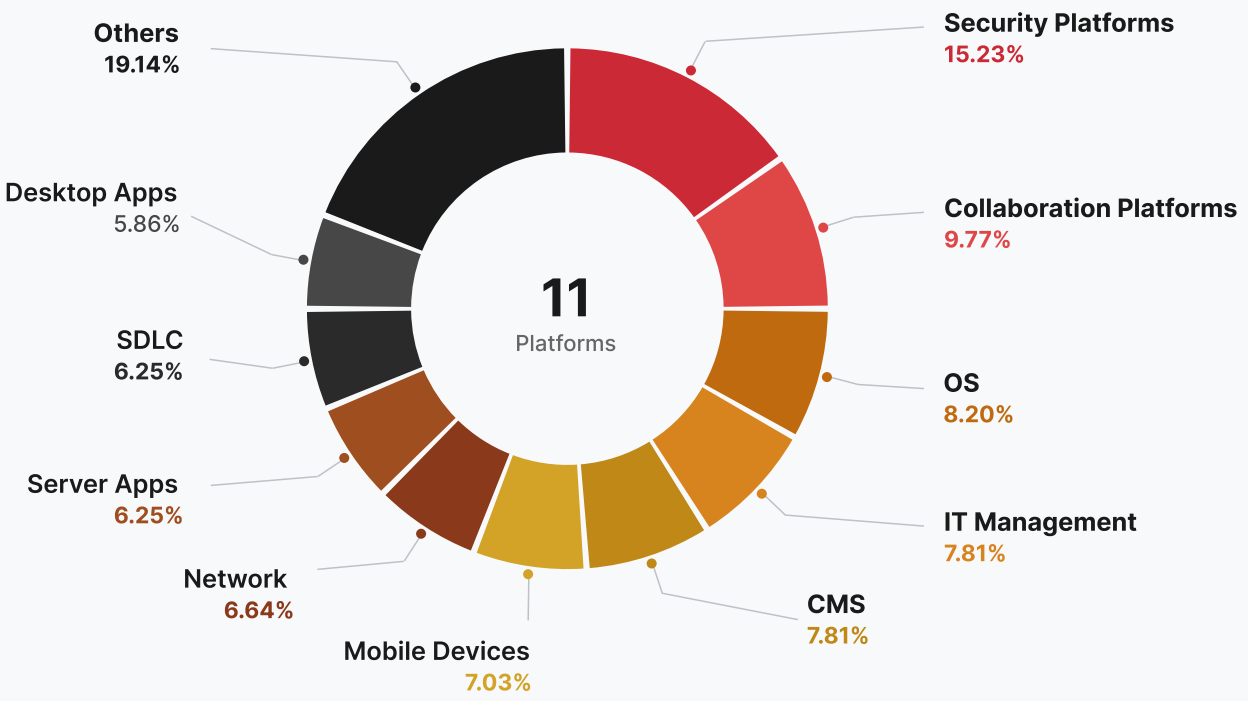

OS · 21 (8.2%)

The OS layer cuts across the entire infrastructure at 8.2% (21 flaws): Windows kernel vulnerabilities (CLFS, NTFS, Win32k, AFD.sys) dominated, with multiple zero-days exploited for privilege escalation in ransomware campaigns.

SDLC Tooling · 16 (6.25%)

The most strategically concerning finding - 6.25% (16 vulnerabilities) from SDLC tooling: CI/CD pipelines, development platforms, and open-source libraries. These aren't high-volume yet, but their blast radius-per-CVE far exceeds any other platform class.

Platform-Level Distribution of Exploited CVEs · 2025*

Edge & Perimeter: first line of defense 35

▼ ▼ ▼

::: SDLC :::

CI/CD Pipelines & Dev platforms

10

Software Libraries

06

::: APPLICATION LAYER :::

CMS	Enterprise apps	Collaboration Platforms	File Transfer
17	11	25	8

▼ HOSTED ON ▼

Server Applications	Database	Data Center Solutions
16	4	5

▼ POWERED BY ▼

Operating System 21

: MANAGEMENT LAYER :

IT management / RMM 20

▼ ▼ ▼

User Zone: endpoints & client-side attack surface 54

Workstations · 36 | Mobile Devices · 18

::: Protecting all zones :::

Security Platforms

39

**Vulnerabilities overlap across the platform stack, a single CVE may be counted in multiple zones.*

Recommendations · Platform-Level Exploitation Defense

RECOMMENDATION #1

Treat Firewalls, VPNs, and EDR as “Crown Jewels”

Treat firewalls, VPNs, and EDR with the **same vulnerability urgency as the other crown jewels they protect**. Adopt **dedicated patching SLAs**, establish **out-of-band update processes**, and enforce **network segmentation** to minimize the blast radius if a security appliance is compromised.

RECOMMENDATION #2

Restrict Management-Console Access

Restrict management console access to a dedicated management network. A large proportion of exploited security and edge device flaws were not exploited through the data plane: they were exploited through admin interfaces that were reachable from the broader network or, worse, the internet. Separating management access eliminates reachability, which eliminates the exploit path.

RECOMMENDATION #3

Supply-Chain Hardening · Treat Build Pipelines Like Production

Implement **software composition analysis**, **pin dependency versions**, **verify build integrity**, and **audit GitHub Actions and pipeline configurations with the same rigor as production code**. SDLC tooling has the highest blast radius per vulnerability of any platform class in the dataset.

CHAPTER 07

The Exploited Flaws : Attack Vector

PRIMARY ATTACK VECTOR

Remote — 142 · 55.5%

Exploitable over the network - targeting internet-facing services, webapps, and exposed management interfaces

↳ Authenticated — 81 · 31.6%

Requires credentials - but credential pools are often large (employees, Saas users, VPN users)

↳ Unauthenticated — 61 · 23.8%

Zero barrier - mass exploitable by botnets and automated scanners. No credentials needed.

Local — 45 · 17.6%

Post-compromise Exploitations that amplify the attack impact - BYOVD, VM escapes, EDR evasion; Privilege Escalation enabling full system compromise.

Phishing — 20 · 7.8%

Requires user interaction - crafted archives, malicious links, and MotW bypasses

Remote Exploitation Local / Post-Compromise Phishing / Social Engineering

OVERLAY DIMENSIONS - RISK MULTIPLIERS

EXPLOIT CHAINS

65 CVEs — 25.4% of all exploited CVEs

Multiple CVEs chained for greater impact. Patching one link breaks the chain: but missing any one gives the attacker the full path.

ZERO-CLICK

10 CVEs — 3.9% of all exploited CVEs

Spyware-grade: no user interaction. Targets mobile platforms (iOS, Android, Samsung) and messaging apps. State-sponsored surveillance.

> *Local privilege escalation vulnerabilities don't make headlines because they can't be the first step. They make breaches catastrophic because they're the second one. They transform a limited foothold into full domain compromise. Teams that deprioritize them because they're not remotely exploitable are misreading the kill chain.*

Paolo Del Mundo

Director of Application Security at The Motley Fool

Recommendations

Controls for the four exploitation patterns observed in 2025: chain-aware prioritization, kernel/driver hardening and email security validation.

RECOMMENDATION #1

For Chained Vulnerabilities

- Start with the "Entry Link":the vulnerability that is most exposed (e.g., internet-facing) or easiest to exploit.
- Followed by, cleaning up of other vulnerabilities in the chain.They can be building blocks of future chains.

RECOMMENDATION #2

Ransomware deployment, EDR kill, domain takeover: all depend on local escalation. Enforce driver blocklists ([4]Microsoft's recommended driver block rules), restrict kernel-mode driver loading, and patch Windows kernel flaws (CLFS, AFD.sys, Win32k).

RECOMMENDATION #3

Conduct email security validation to assess whether your email security solution can effectively detect and block phishing and BEC attacks.

CHAPTER 08

Exploitation Objectives

The analysis of immediate attacker objectives across 256 actively exploited CVEs reveals a stark 2:1 ratio - **for every CVE that steals data, bypasses authentication, or escalates privileges, there are two that run arbitrary commands**, making code execution the dominant attacker outcome by a wide margin.

EXPLOITATION OBJECTIVES · 256 CVES IN 2025 · N=256

INTELLIGENCE HARVESTING

42 16.4%

Data Disclosure and Credential Access combine to give attackers visibility into target environments. Credential theft **carries disproportionate downstream impact**: a single stolen NTLM hash or API key can unlock entire environments.

INITIAL ACCESS

35 13.7%

The **"door openers"**: authentication bypasses and default credentials that don't execute code themselves but **hand the attacker a valid identity.** Many of the most impactful exploitation chains pair an Initial Access flaw with a code execution vulnerability.

CODE EXECUTION

129 50.4%

Half of all exploited vulnerabilities gave attackers the **ability to run arbitrary commands:** making code execution the dominant attacker outcome by a wide margin (a stark 2:1 ratio over every other objective).

PRIVILEGE ESCALATION

31 12.1%

The amplification layer - never the first step, always the second. Windows kernel and driver flaws dominate, with CLFS, DWM, Hyper-V, and AFD.sys each contributing multiple exploited CVEs.

DEFENSE EVASION

14 5.5%

Small in count but outsized in strategic importance. Every BYOVD EDR kill, sandbox escape, and MotW bypass exists to make another vulnerability in this dataset more lethal.

OTHER / UNCATEGORIZED

5 ~2%

DoS, integrity tampering, and Lateral Movement: the residual after the five primary buckets. Smaller share, occasionally critical to specific environments.

Recommendations

RECOMMENDATION #1

Application Behavioral Controls

Half of all exploitation ends in arbitrary code execution. **Deploy application behavioral whitelisting and process execution controls** on critical servers. If your web server, database, or edge appliance spawns cmd.exe, PowerShell, or bash unexpectedly: that's not a false positive, that's the 50.4% happening in real time.

RECOMMENDATION #2

Credential Hygiene

Rotate credentials aggressively on exposed services, **enforce MFA wherever possible**, monitor for credential use from unexpected locations or devices, and treat any exposed heap dump, configuration file, or debug endpoint as a credential breach: because attackers do.

CHAPTER 09

Conclusion : What the Data Says Altogether?

Individually, each section of this report tells its own story. Together, they reveal something the numbers alone don't show: a **system-level failure** in how the industry approaches vulnerability management.

Security Products Are Now the Most-Exploited Category

Firewalls, VPNs, EDR, and IAM **account for the single largest share of exploited flaws at 15.2%**. The tools deployed to defend the network are the network's most exploited attack surface. The attack vector data explains why: these products sit at the perimeter, exposed to the internet by design. **When 23.8% of all exploited flaws require no authentication, and your firewall's management console is one of them, the device meant to be your front door becomes the attacker's front door.**

The Objectives Reveal a Kill Chain, Not Isolated Techniques

Half of all exploited CVEs delivered code execution: but the other half didn't need to. 13.7% gave attackers a valid identity through authentication bypass alone. Another 12.1% elevated privileges. 5.5% existed solely to disable the security controls meant to catch everything else. An Initial Access flaw bypasses the VPN, a Privilege Escalation flaw owns the domain, a Defense Evasion flaw kills the EDR, and a code execution flaw deploys ransomware: four CVEs from four different sections of this report, working as a single operation.

Threat Actors Are Sharing Tools and Targets

25.4% of exploited CVEs were chained together in multi-vulnerability sequences. SAP NetWeaver (CVE-2025-31324) drew 9 distinct threat actors: six state-sponsored and three criminals: onto the same target. WinRAR (CVE-2025-8088) was weaponized by all four major nation-state groups simultaneously. These aren't theoretical chains from a red team exercise; they are observed campaigns where the attacker already assembled the pieces.

Zero-Days and N-Days Are Mixed in the Same Operation

40.6% of exploited flaws had no patch available when exploitation began. The platform with the most zero-days: Microsoft, with 23, is also the platform where most privilege escalation and defense evasion flaws sit. Attackers aren't choosing between zero-days and known vulnerabilities; they are mixing them in the same operation, using a zero-day to get in and an unpatched n-day to move laterally.

The Convergence Is What Makes This Moment Different

It's not any single trend: it's the convergence. **The noise is growing, the exploitation window is collapsing, the attacker ecosystem is sharing tools and targets, and AI is compressing every timeline further.** The recommendations that follow aren't aspirational: they're the minimum response to a landscape where the old model of scan, score, and patch in order has already failed.

THE SHIFT TO MAKE

Vulnerability management has been a patch-tracking discipline for two decades. The 2025 data makes the argument that it should become a threat-prioritization discipline instead: one where the question moves from "what's left to patch" to "what is being weaponized right now, and how exposed are we." Every chapter in this report points the same direction.

CHAPTER 10

Your 2026 Survival Playbook

Based on the threat intelligence presented in this report, HivePro recommends the following prioritized actions to reduce organizational risk exposure.

Top 10 Prioritized Recommendations

#	ACTION	ADDRESSES	PRIORITY	TIMELINE
1	**Threat-informed vulnerability prioritization**	Optimizes Patch Cycle, removes patch noise	CRITICAL	Immediate
2	**Adversarial Exposure Validation**	Validate Defenses against real-world threats	CRITICAL	Immediate
3	**AI-powered email & endpoint defense**	7.8% Vulnerability Exploitation	CRITICAL	Immediate
4	**Zero Trust architecture enforcement**	Reduces Blast Radius	CRITICAL	0–6 months
5	**Ransomware response playbook creation and testing**	12.9% CVE Exploitation yields ransomware	HIGH	0–3 months
6	**Supply chain security program**	Growing Attack Vector	HIGH	3–6 months
7	**Network segmentation / micro-segmentation**	Lateral movement (27% Network AV)	HIGH	3–9 months
8	**Dark web & leak site monitoring**	Keeps 31.6% exploitation activities in check	HIGH	3–6 months
9	**AI governance & model security**	AI-targeted attacks, shadow AI	MEDIUM	3–6 months
10	**Quantum-safe crypto transition plan**	Harvest-now-decrypt-later threat	MEDIUM	6–18 months

Network segmentation is one of the highly underrated control that changes the math on both zero-days and chained exploits. It doesn't prevent initial access but it constrains what an attacker can reach once they're in. Most organizations have some segmentation on paper. Very few have validated that it actually holds against the lateral movement techniques in active use. There's a meaningful difference between a segmentation design and a segmentation capability.

Matt Walker

Cybersecurity Leader, Technologist

Hive Pro

2026 Threat Predictions

1 AI Arms Race Escalation

Offensive AI will become a default, not the exception. Expect **autonomous attack chains that adapt in real-time to defensive responses**.

2 Vulnerability Acceleration

CVE volume will exceed 70,000. Time-to-exploit will drop further. **Threat-informed prioritization becomes non-negotiable.**

3 Supply Chain Systemic Risk

A **major upstream software compromise will impact thousands of downstream organizations simultaneously**. Vendor security will become a board-level conversation.

1% IS ALL THAT MATTERS

In 2025, 131 CVEs were published every day. Less than 1 per day was exploited in the wild. 104 were zero-days: exploited before a patch even existed.

Your patching strategy is either threat-informed, or it's guesswork at scale.

Want to dive deeper into your threat landscape?
Connect to a Threat Intelligence expert today!

Connect with a Threat Intelligence Expert

Scan the QR to
Subscribe Our Threat Digest