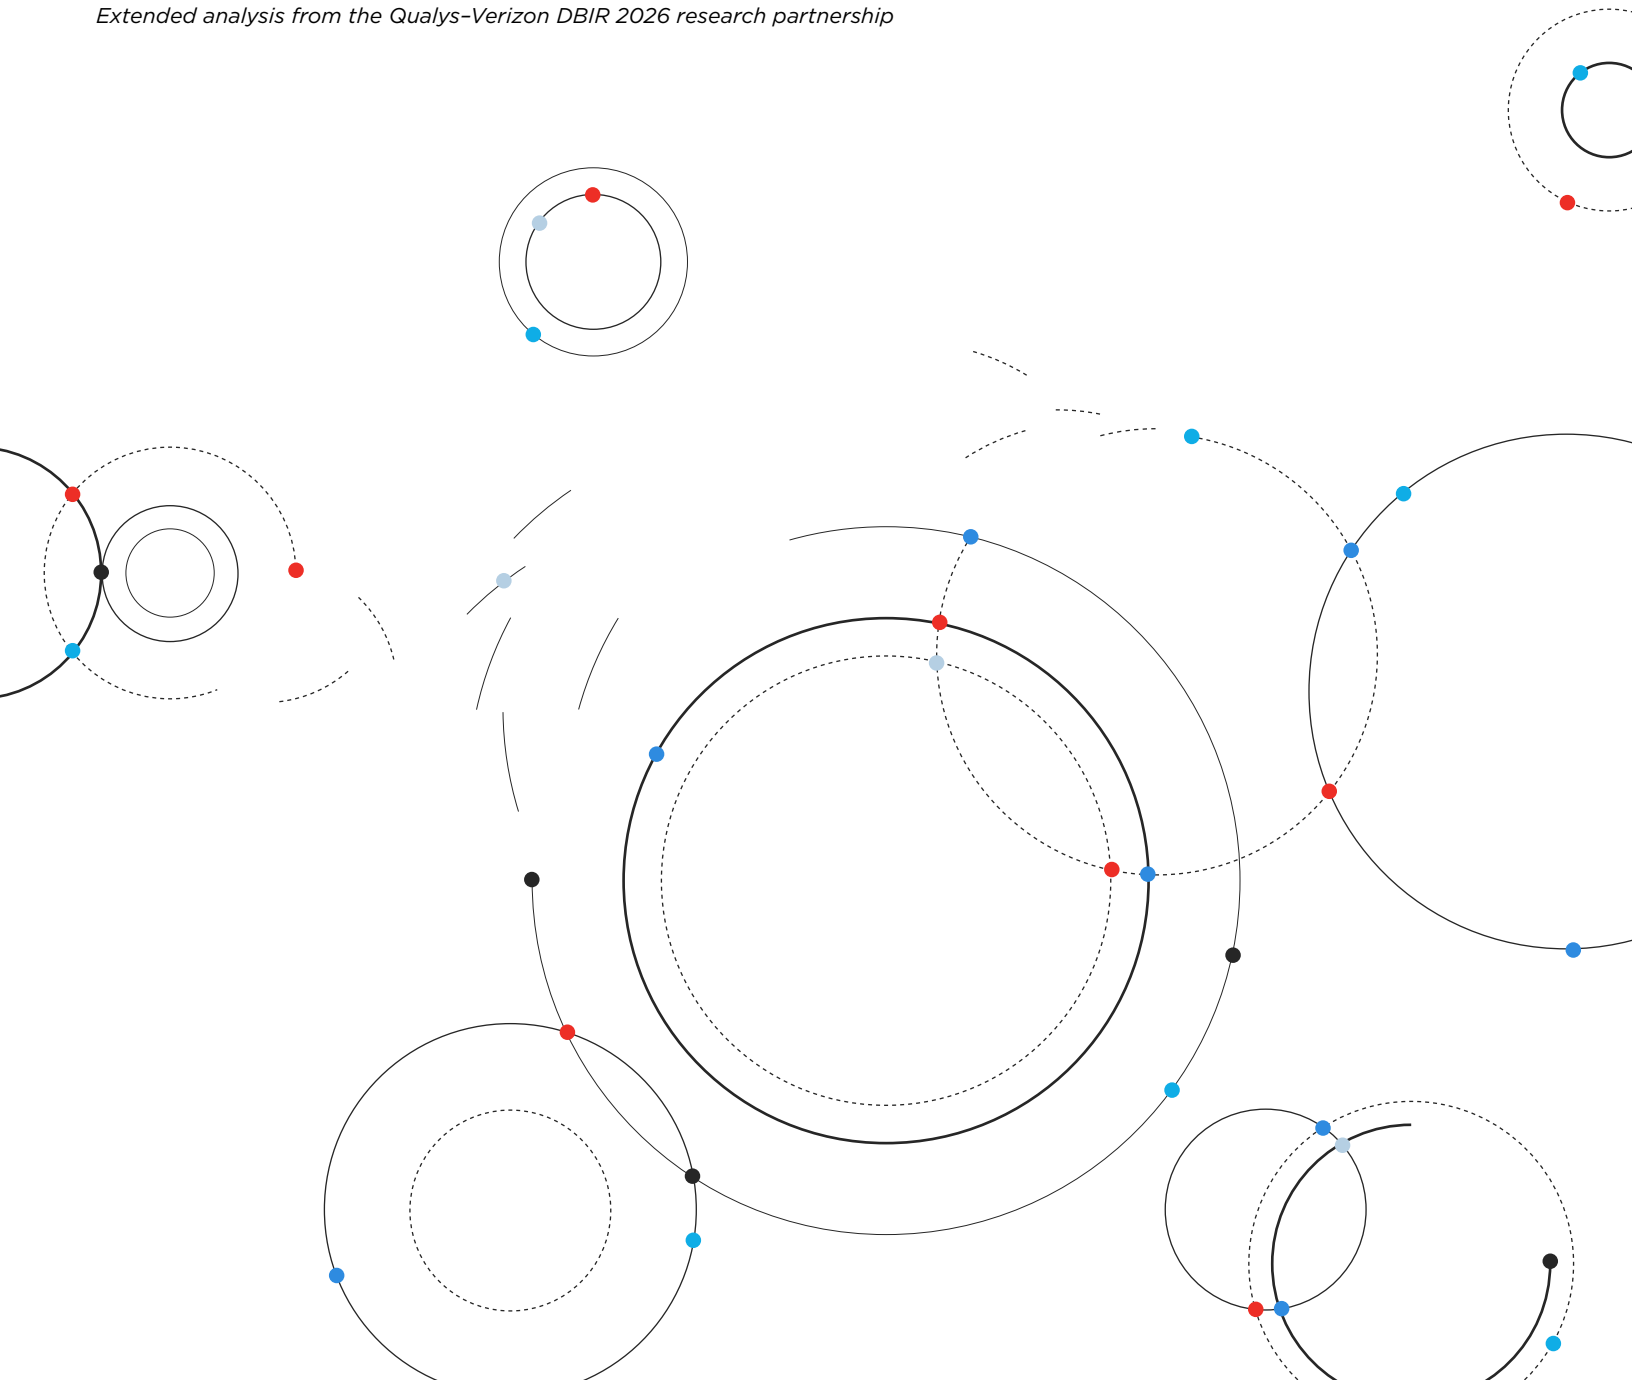


The Broken Physics of Remediation

The future of risk management is autonomous.

Extended analysis from the Qualys-Verizon DBIR 2026 research partnership



Foreword

The threat landscape has fundamentally shifted, and the defenses most organizations rely on were not built for what comes next. As I meet with customers and hear firsthand about the challenges they're facing, the threats I see are quieter and more structural: attack surfaces expanded beyond what teams can govern, identity sprawl that outpaces policy, and remediation workflows still built on manual execution. These, not headline-grabbing ransomware strains or sophisticated zero-days, are what will define enterprise risk. In an era where adversaries increasingly operate at machine speed, any architecture that depends on human-speed response carries structural risk.

The data in this report validates that reality across more than one billion CISA KEV remediation records spanning 10,000 organizations over four years. The average Time-to-Exploit has collapsed to negative seven days with adversaries weaponizing vulnerabilities before patches even exist. Critical vulnerability volume has surged 6.5x, yet the percentage still open at Day 7 and Day 30 has worsened. Of the 52 high-profile weaponized vulnerabilities we tracked with complete exploitation

timelines, 88% were remediated slower than they were exploited. Half were weaponized before public disclosure. The traditional patch race, the belief that reducing Mean Time to Remediate can outrun the attacker, has hit a hard mathematical ceiling. Meanwhile, less than one percent of all disclosed CVEs represent confirmed, weaponized, remotely exploitable risk. Organizations are burning remediation cycles on theoretical exposure while genuinely exploitable gaps persist.

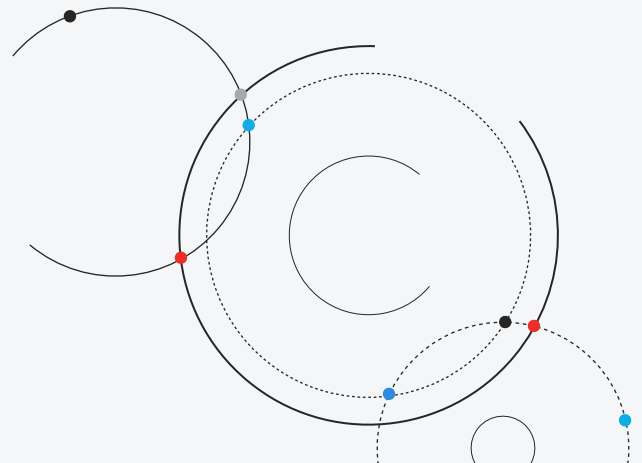
The mandate is clear. We must match autonomous offense with autonomous defense. This requires a foundational architectural shift away from reactive human triage and toward a Risk Operations Center (ROC) that fuses embedded intelligence, deterministic confirmation of actual exploitability, and autonomous remediation into a single operational loop. The goal is not to remove human judgment, but to elevate it, shifting practitioners from executing tactical actions to governing the policies that guide autonomous systems.


Welcome to the new operational reality.



A handwritten signature in black ink that reads "Sumedh Thakar".

Sumedh Thakar
President and CEO
Qualys





Introduction

Before you read the data in this report, it helps to understand the adversary it measures.

The attacker's timeline is the only one that matters — and that timeline is predictable. Adversaries do not innovate; they repeat what works. Once a threat actor invests in exploiting a complex product, they do not stop at one CVE. They systematically exploit that entire software family until defenders are overwhelmed. The remediation data in this report reflects this: the same product classes appear in the long tail of unresolved exposure, not because they are unknown, but because the operational model cannot close them at speed.

The path of least resistance has shifted from the endpoint to the edge and from there, deeper into the enterprise software organizations implicitly trust. Highly privileged, deeply embedded, and rarely owned by security teams, these systems offer adversaries the most efficient return on a single exploit. The remediation data shows why: these are the asset classes where the Manual Tax is heaviest, the

exposure window is longest, and the confirmation gap this report identifies is most acute.

The consequences are predictable. The remediation failures documented in this report trace a direct line to ransomware outcomes. But ransomware is the symptom. The disease is the cumulative, unresolved exposure that gave the attacker a window to act — what this report measures as Risk Mass.

Every generation of cybersecurity emerged in response to a platform shift. New technology created new risk, and defenders adapted. That cycle was linear and survivable. What is emerging now is not another platform shift. It is the first time the adversary itself is becoming autonomous. Now, offensive agents can discover, weaponize, and execute faster than any human-staffed operation can respond. The defensive side must make the same transition — and this report measures the cost of every day the transition is delayed.

What follows is the data that proves it.



Saeed Abbasi

Head of Threat Research Unit
Qualys

A Note on This Edition

This report was originally published in March 2026. Following its citation in the Verizon 2026 Data Breach Investigations Report, we have added Section 7—an appendix presenting extended analysis developed through our research partnership with the Verizon DBIR team, which could not fit within the published DBIR report.

Table of Contents

- 5 The End of Human-Scale Remediation
- 9 The Broken Physics of Remediation
- 13 The Physics Gap—Attacker Speed vs. Defender Speed
- 21 Risk Mass—From Counting Vulnerabilities to Measuring Exposure
- 25 The Filter—From Prioritization to Confirmation
- 31 Operationalize or Fail
- 35 Appendix: Survival of the Vulnerable—Extended Analysis from the DBIR 2026 Partnership

Executive Summary


According to Google Threat Intelligence Group (GTIG), as published in the [Mandiant M-Trends 2026 report](#), the mean Time-to-Exploit (TTE) collapsed to -1 day in 2024, and then to -7 days in 2025. Adversaries are weaponizing vulnerabilities and compromising systems before a patch is released or the moment it is disclosed. This represents a fundamental breakage in the operational model of remediation.

For the last decade, the industry operated on the assumption of a patch race—that if we could just reduce MTTR, we could outrun the adversary. You cannot solve a “minutes” problem with a “months” solution.

This report draws on over 1 billion CISA KEV remediation records across 10,000+ organizations (2022–2025) to quantify the widening gap. We show that KEV vulnerability volume has grown 6.5x while the percentage of critical vulnerabilities still open at Day 7 has increased—proof that manual remediation has hit a hard ceiling. We introduce Average Window of Exposure (AWE) and Risk Mass as complements for MTTR, metrics that capture not just response speed but cumulative exposure at scale.

The findings go further. Even among vulnerabilities scored as highly exploitable, confirmation testing reveals that only a fraction pose validated risk once compensating controls are accounted for—meaning organizations are spending remediation cycles on threats already mitigated, while genuinely exploitable gaps persist.

As the velocity of exploitation accelerates—increasingly driven by adversaries leveraging AI to automate reconnaissance, weaponization, and attack execution at machine speed—human intervention becomes not merely slow, but structurally incompatible as the primary line of defense. The path forward is clear: end-to-end operationalized remediation through a Risk Operations Center (ROC) approach—automating the full cycle from detection through validation to action. Risk whack-a-mole does not scale. Repeatable, machine-speed operationalization is the only response the threat environment still permits.



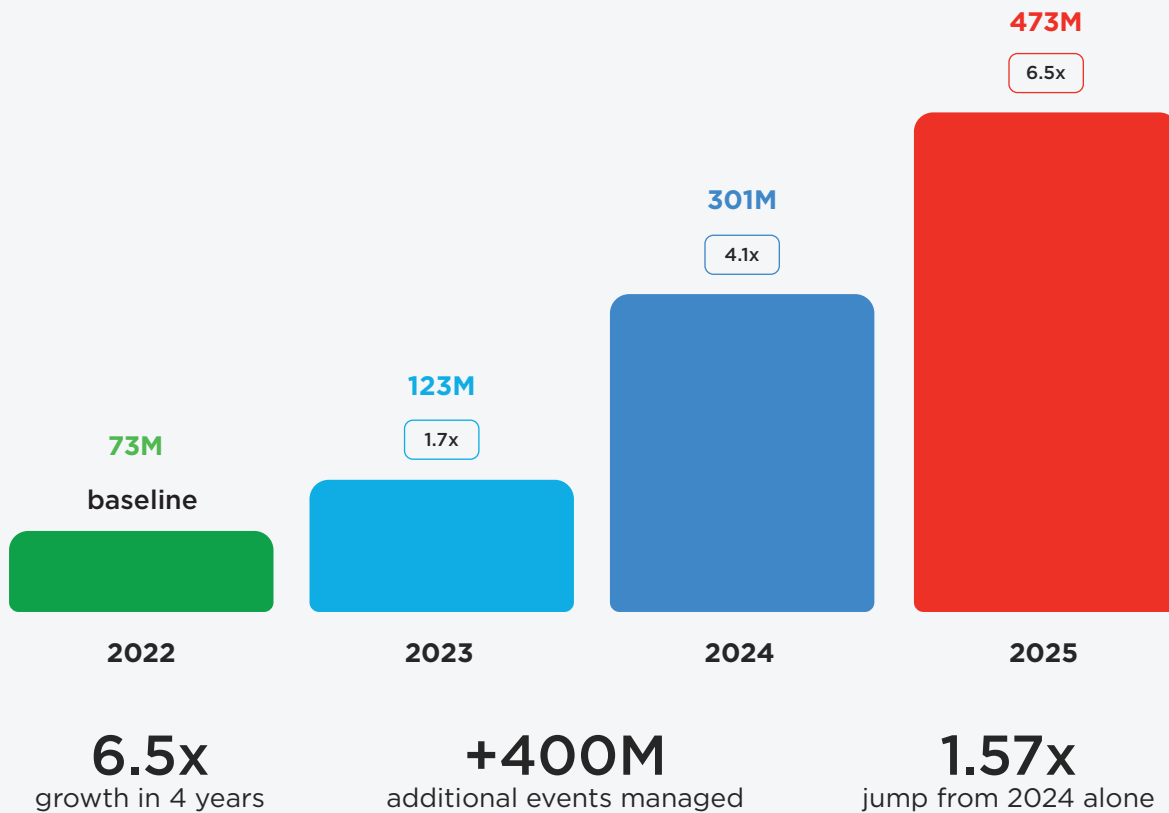
1

The End of Human-Scale Remediation

Manual remediation was designed for a world that no longer exists — one where vulnerability volumes grew linearly, exploit timelines were measured in weeks, and a well-staffed team could reasonably expect

to triage, ticket, and patch its way to safety. That world ended quietly, somewhere between the 73 millionth and the 473 millionth remediation event.

The CISA KEV Volume Explosion




Vulnerability events closed per year across 10,000+ organizations


The Human Ceiling

We analyzed over 1 billion CISA KEV vulnerability remediation records across 10,000+ organizations from 2022 through 2025. The dataset represents one of the largest longitudinal studies of enterprise remediation behavior ever conducted. Each record represents a single vulnerability instance detected on a single asset, tracked from detection through remediation

closure. It reveals a structural limit that no amount of staffing, process maturity, or executive urgency can overcome.

The volume of closed vulnerability events grew 6.5x in four years — from approximately 73 million in 2022 to 473 million in 2025 — driven by both rising vulnerability volume and broader organizational coverage.





The volume of closed vulnerability events grew **6.5x** in four years — from approximately **73 million** in 2022 to **473 million** in 2025 — driven by both rising vulnerability volume and broader organizational coverage.

Even controlling for that growth, the percentage of critical CISA KEV vulnerabilities still open at Day 7 did not improve. It got worse — rising from 56% in 2022 to 63% in 2025.

This is the human ceiling. Teams worked harder and closed more tickets in absolute terms, but the rate of incoming risk

outpaced the rate of remediation. The backlog didn't shrink; it compounded. You cannot scale a workforce 6.5x to match a 6.5x increase in exploitable risk — and even if you could, the data shows it would not be enough. The constraint is not effort. It is the operational model itself.


The Intelligence Imperative

The collapse of remediation timelines carries a second-order consequence that is less discussed but equally urgent: it fundamentally changes how threat intelligence must be consumed for operational remediation decisions.

Strategic threat intelligence — attribution analysis, TTP mapping, actor motivation, campaign tracking — remains essential to the security program. It informs defensive architecture, shapes investment priorities, guides threat hunting, and provides the contextual understanding that no automated system can replicate.

What has changed is the operational layer — the point where intelligence must translate into a specific remediation action on a specific asset. When remediation was a human-driven process, tactical intelligence could afford to travel through human channels — an analyst reading a report, interpreting its relevance, and routing it into a workflow. That model assumed time existed between insight and action. At a Time-to-Exploit of -7 days, that assumption no longer holds for operational remediation decisions.

For this tactical layer, intelligence must be embedded directly into the platforms where remediation decisions are made. It must arrive with environment-specific context — not just “this CVE is being exploited in the wild,” but “this CVE is being exploited, it exists in your environment, your current controls do not mitigate it, and here is the action to take.” It must be structured enough for a machine to parse and specific enough to drive an autonomous decision. Strategic intelligence informs the program. Operational intelligence, delivered at machine speed, informs the action. Both are necessary. But for the remediation pipeline specifically, human-speed routing is no longer fast enough.



2

The Broken Physics of Remediation

MTTR — Mean Time to Remediate — has earned its place as the industry’s default remediation metric for good reason. It measures something real and operationally important: how quickly security and IT teams respond once a vulnerability enters the remediation pipeline. As a measure of operational efficiency, MTTR remains valuable — it tells you whether your teams are getting faster, whether SLAs are being met, and whether process improvements are working.

What MTTR was not designed to measure is business risk exposure. It does not capture

what happened before the first ticket was opened, how long the environment was exposed while remediation was underway, or the compounding risk absorbed by every asset that remained unpatched while the average was being calculated. MTTR measures the speed of the response. It does not measure the duration of the exposure. As exploit timelines compress toward and past zero, that distinction becomes operationally critical — not because MTTR is the wrong metric, but because it is an incomplete one for the risk question that boards and executives increasingly need answered.



AWE captures what MTTR ignores: the full duration between the moment a vulnerability becomes weaponized and the moment it is remediated across the environment.

Average Window of Exposure

AWE captures what MTTR ignores: the full duration between the moment a vulnerability becomes weaponized and the moment it is remediated across the environment. Think of it as measuring the entire race — not just the final sprint. If the attacker arrives on Day 1 and you finish patching on Day 21, your AWE is 21 days. If MTTR told you “8 days,” it averaged the fastest closures with the slowest — producing a single number that obscured the fact that a third of your assets were still exposed weeks later.

To address a structural blind spot in how the industry quantifies remediation effectiveness, Qualys President and CEO Sumedh Thakar developed Average Window of Exposure (AWE). Where MTTR measures operational velocity — how fast teams close tickets — AWE measures

strategic exposure across the full asset population. The distinction is not academic. As exploit timelines compress past zero-day thresholds, the gap between “how fast we responded” and “how long we were exposed” becomes the difference between a metric that rewards process and one that reflects risk. AWE bridges that gap — capturing the full window from exploitability to remediation closure at scale and surfacing the long-tail exposure that averaging inherently conceals.

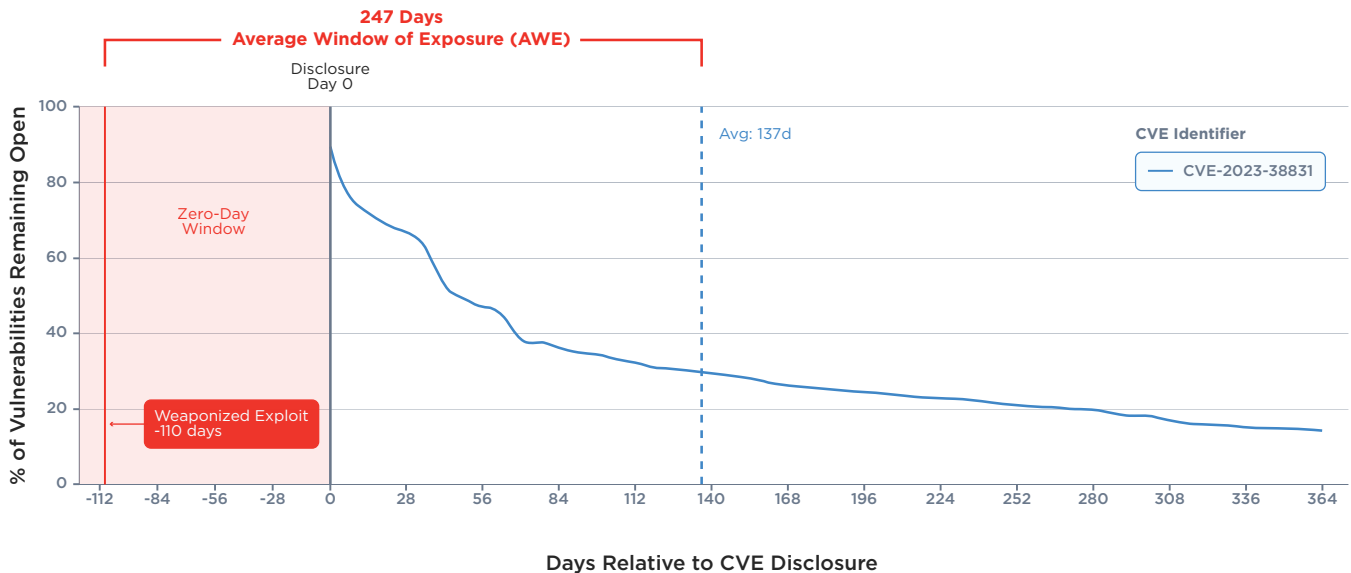
Case in Point: WinRAR Code Execution (CVE-2023-38831)

Consider CVE-2023-38831, the WinRAR Code Execution vulnerability, as the principle made concrete. Weaponized exploits were circulating 110 days before public disclosure, a window in which defenders had no signal, no patch, and

no ticket to open. Once the CVE was published, the average organization required an additional 137 days to remediate across its environment. The result is an Average Window of Exposure of 247 days: more than eight months during which a confirmed, actively exploited vulnerability sat open across the asset population. MTTR, measured only from

disclosure forward, captures less than half of that story. Nor does the exposure end at the average. A full year after disclosure, roughly 15% of vulnerable systems remained unpatched, a structural long tail invisible to any metric anchored to mean closure time. This is the gap AWE makes visible, and the gap that no acceleration of human-paced remediation can fully close.

AWE in Practice for WinRAR Code Execution (CVE-2023-38831): 247 Days of Exposure



GTIG’s 2024 analysis of 112 exploited-in-the-wild vulnerabilities found that the average Time-to-Exploit collapsed to -1 day — meaning that for vulnerabilities confirmed as actively exploited, weaponization now occurs, on average, *before* a patch exists.

The defender’s clock tells a different story. Our survival curve analysis across CISA KEV critical vulnerabilities in 2025 shows that at the moment of disclosure, 85% of vulnerable assets remain unpatched. At one week, 63%. At the average remediation mark of approximately 21 days, one in three

assets — 33% — is still open. And at 90 days, nearly 12% of the attack surface persists.

It is worth noting that the 15% patched before KEV addition represents organizations that have moved beyond reactive remediation — applying risk-based prioritization through scoring systems such as TruRisk and advanced remediation capabilities to address likely-exploitable vulnerabilities before they are confirmed as actively exploited. They are the proof that compressing the curve is possible. They are also the minority.

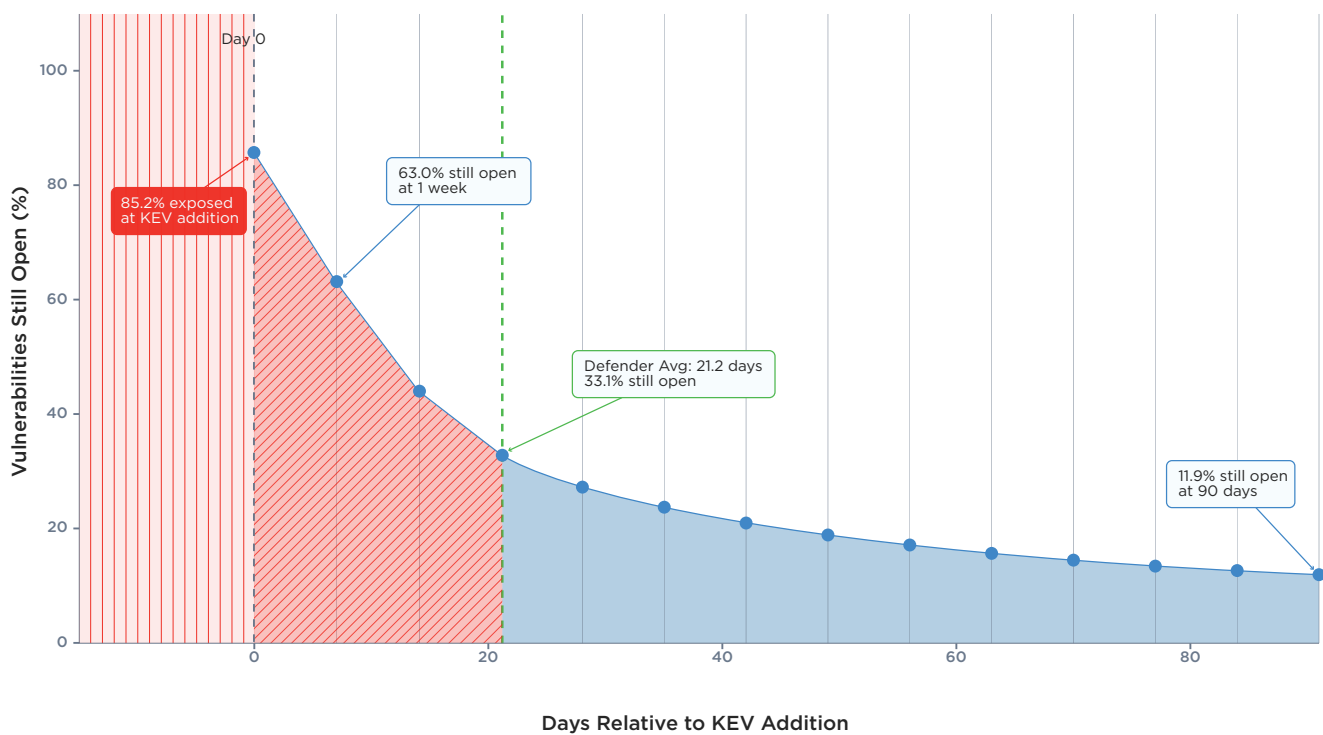
The Breach Window

The shaded area beneath the survival curve is not an abstraction. It represents the cumulative exposure your organization absorbs for every day remediation remains incomplete — what we define as Risk Mass, explored in depth in Section 4 of this report. Every point on that curve is an open door. The steeper the initial drop, the more effective your first-response automation. The longer the tail, the greater the accumulated risk that no dashboard will show you.

Defenders begin remediating at Day 0. For the majority of critical zero-day vulnerabilities, attackers were already there. The gap between those two starting lines — and the slow decay of the remediation tail — is where breaches live. No amount of manual acceleration closes it. Only automated, operationalized remediation can compress the curve fast enough to change the outcome.

The Breach Window

Attacker exploitation speed vs. defender remediation — the gap where breaches happen



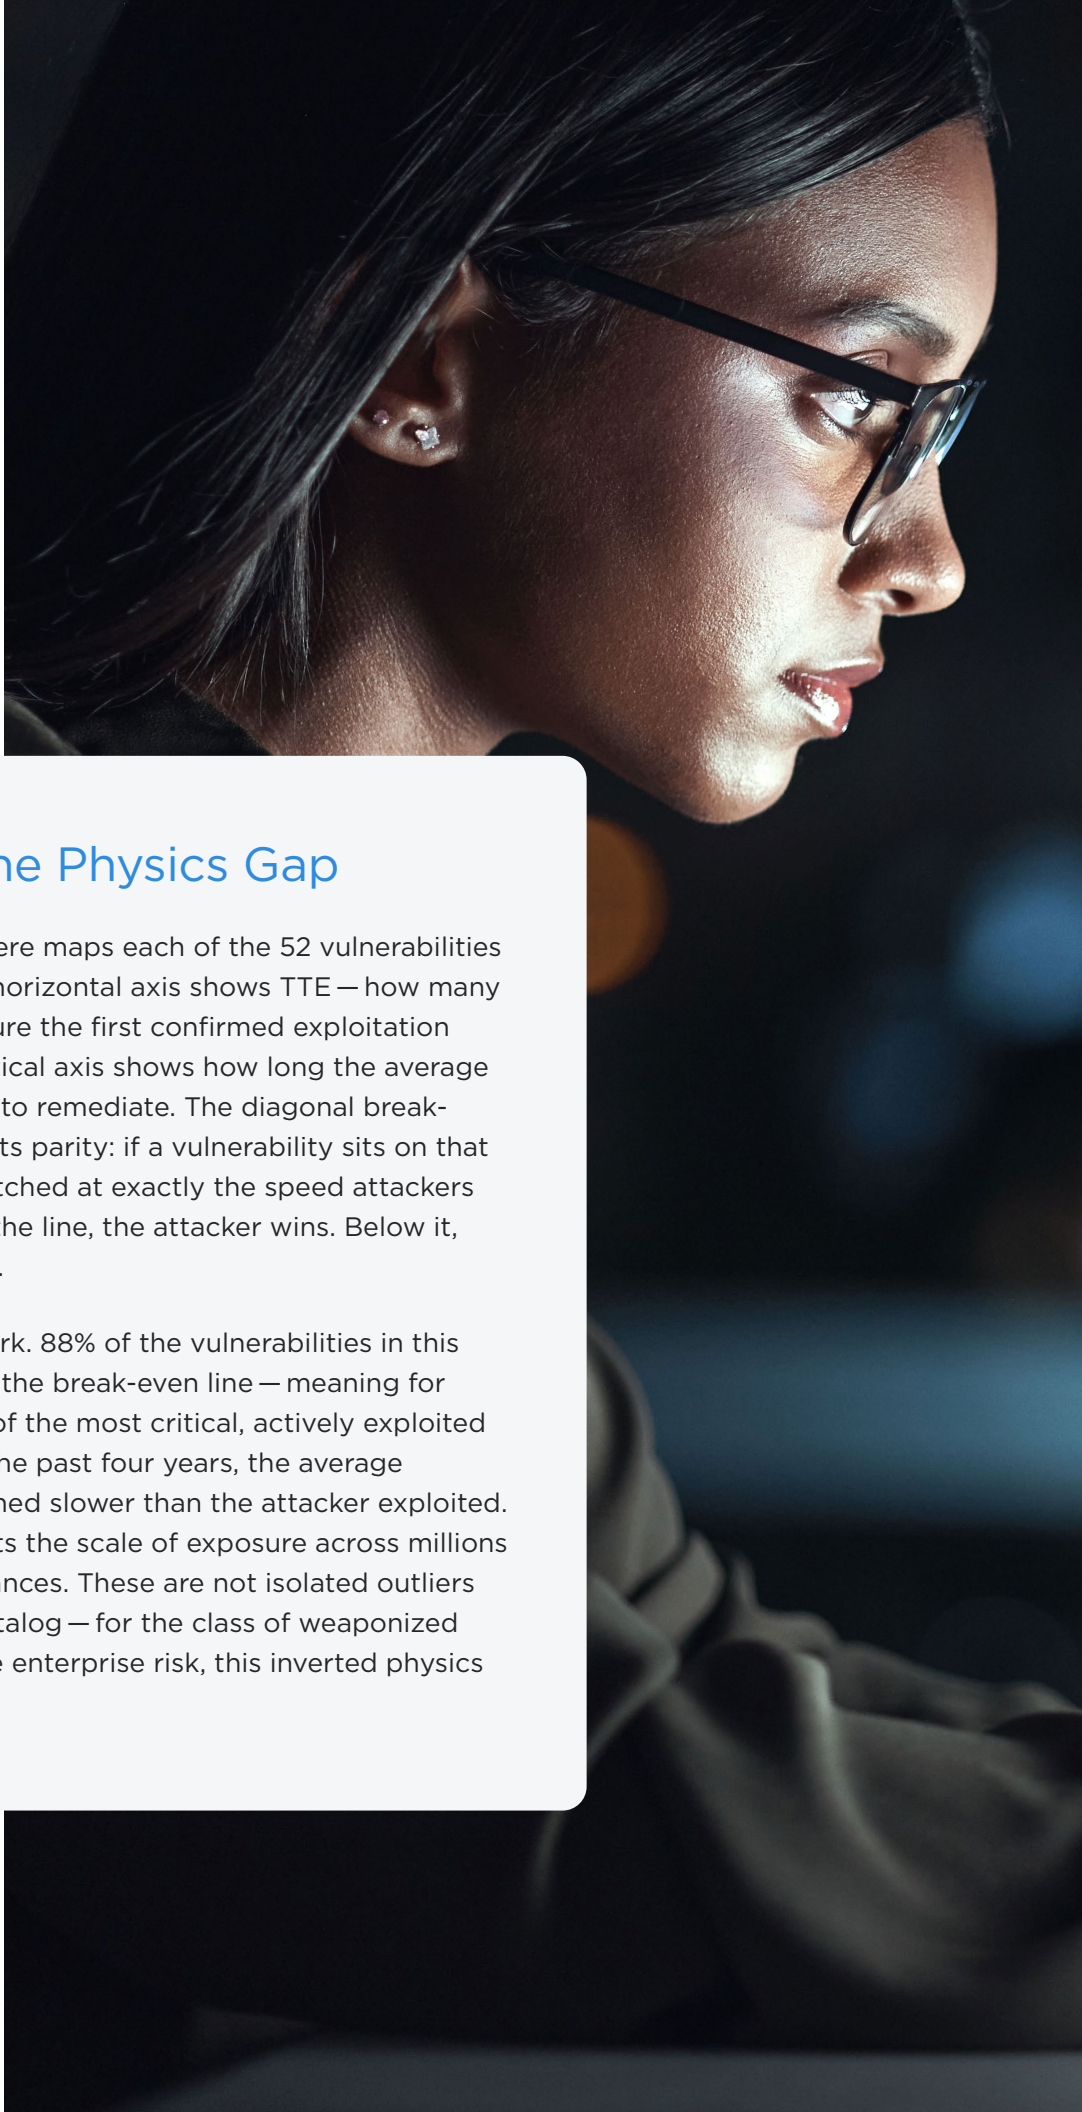



3

The Physics Gap — Attacker Speed vs. Defender Speed

The previous section established that the average enterprise leaves critical vulnerabilities exposed for 21 days while attackers exploit in less than one. This section examines what that asymmetry looks like across 52 individual CISA KEV vulnerabilities for which our threat intelligence — corroborated across multiple commercial and proprietary intelligence sources — maintained complete exploitation timelines with confirmed first-


exploitation dates. Because this cohort is defined by intelligence completeness rather than random sampling, it may skew toward higher-profile vulnerabilities that attracted broader research attention. Even with that caveat, the gap is worse than aggregate statistics suggest. Half of these vulnerabilities were exploited before public disclosure. For the remainder, defenders still lost the majority of engagements.



Reading the Physics Gap

The scatter plot here maps each of the 52 vulnerabilities on two axes. The horizontal axis shows TTE — how many days after disclosure the first confirmed exploitation occurred. The vertical axis shows how long the average organization took to remediate. The diagonal break-even line represents parity: if a vulnerability sits on that line, defenders patched at exactly the speed attackers exploited. Above the line, the attacker wins. Below it, the defender wins.

The results are stark. 88% of the vulnerabilities in this dataset fall above the break-even line — meaning for the vast majority of the most critical, actively exploited vulnerabilities of the past four years, the average organization patched slower than the attacker exploited. Bubble size reflects the scale of exposure across millions of vulnerable instances. These are not isolated outliers within the KEV catalog — for the class of weaponized threats that define enterprise risk, this inverted physics is the norm.



The Zero-Day Zone

Twenty-six of the 52 vulnerabilities — exactly half — have negative TTE, meaning exploitation has happened before the CVE was publicly disclosed. An additional three were exploited on the day of disclosure itself. The “Day 0” starting line that remediation programs are built around does not represent the beginning of the race for the majority of these critical threats. It represents the point at which defenders discover they are already behind.

The scale of pre-disclosure exploitation varies widely. Win Kernel EoP (CVE-2024-21338) was exploited 182 days before disclosure. WinRAR Code Execution (CVE-2023-38831) was weaponized 110 days early. ProxyNotShell (CVE-2022-41040) and its companion vulnerability were exploited 85 days before the advisory. FortiOS Auth (CVE-2024-55591) was weaponized 60 days before publication. Follina (CVE-2022-30190) was exploited a full month before disclosure. Others — SmartScreen, Oracle EBS, Win CLFS EoP, MSC EvilTwin, Sophos Auth, Citrix RCE, PAN-OS CmdInj, PAN-OS Auth, ActiveMQ, Cisco IOS XE — were all exploited days to weeks before defenders had any public knowledge of their existence.

What makes the zero-day zone particularly damaging is not just the head start — it is what happens after the patch arrives. Cisco IOS XE was exploited 30 days before Day 0; average remediation was 263 days. ActiveMQ RCE was exploited 15 days early; average close was 459 days. Spring4Shell was exploited 2 days before disclosure; average close was 266 days. The attacker’s advantage was measured in days or weeks. The defender’s response was measured in months or years.



The Manual Tax

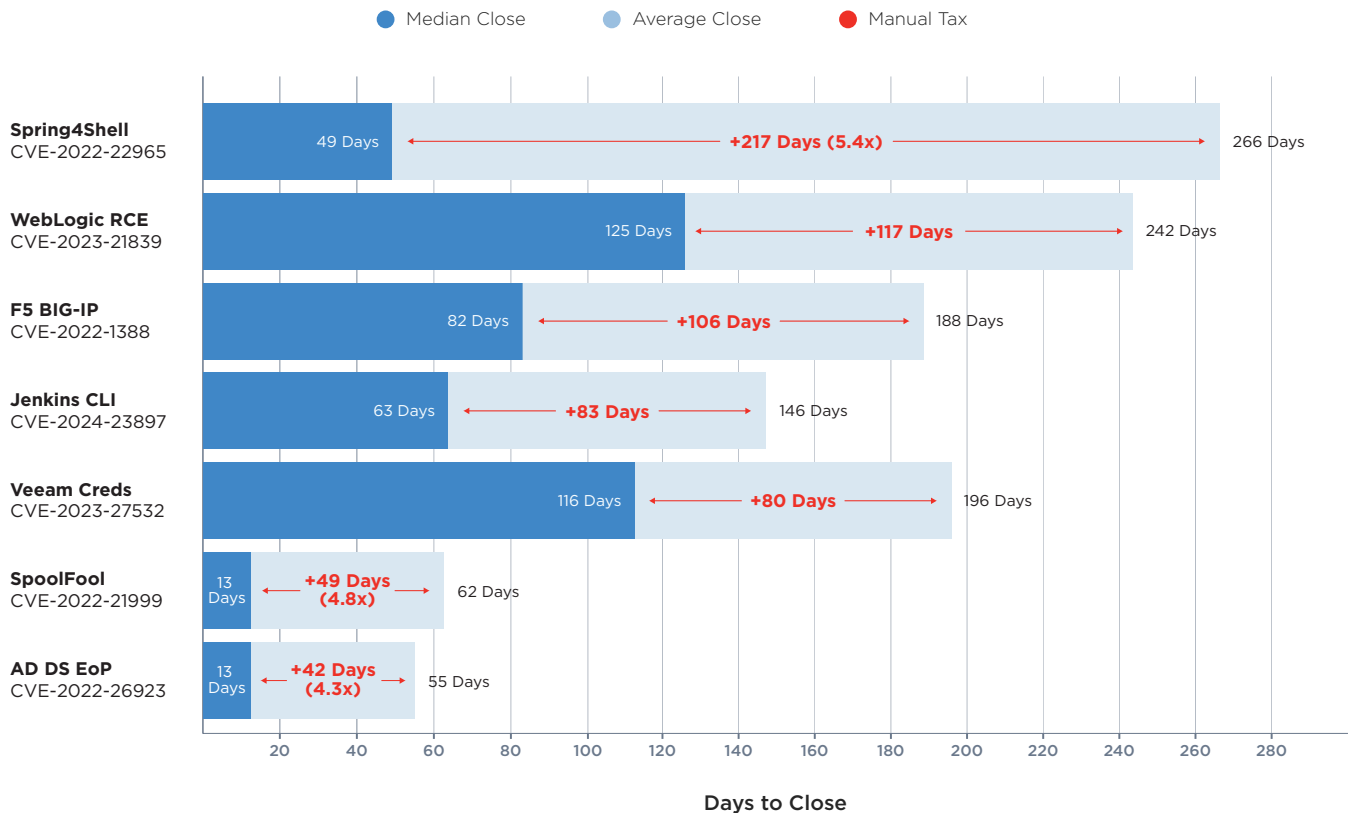
The gap between median and average remediation time reveals a consistent pattern across the dataset that we term the Manual Tax — the cost imposed by the long tail of assets that manual processes cannot reach at speed.

The Manual Tax is not a flat number — it is a multiplier. For Spring4Shell, the average is 5.4 times the median. For SpoolFool, 4.8 times. For AD DS Elevation of Privilege, 4.3 times. The bottom half of the distribution — forgotten servers, remote laptops, shadow IT, assets deferred to the next sprint — takes four to five times longer than the

organization's best performers. While the first 50% of tickets close at a pace that looks defensible on a dashboard, the remaining tail drags actual exposure out by three to nine months.

For some vulnerabilities, even the median is a loss. Cisco IOS XE (TTE: -30 days, median: 232 days), ActiveMQ RCE (TTE: -15 days, median: 530 days), F5 AuthBypass (TTE: 4 days, median: 231 days). For these, even the fastest half of organizations took seven to seventeen months to close vulnerabilities that were weaponized weeks or months before disclosure.

The Manual Tax: Median vs. Average Days to Close



The Infrastructure Divide

The Manual Tax is not evenly distributed. When we segment the 52 vulnerabilities by asset type, three distinct remediation realities emerge — each requiring a different operational approach.


Endpoint and client-side assets patch the fastest. Median close times consistently fall under 14 days, even at volumes exceeding 12 million events per vulnerability. Automated patching infrastructure exists for these assets and it delivers. The operational priority for endpoints is extending that speed advantage with pre-disclosure detection and mitigation capabilities, so that the growing number of vulnerabilities exploited before disclosure can be addressed with the same efficiency.

Edge and perimeter devices — firewalls, VPNs, gateways — carry the highest strategic risk per vulnerability. They are internet-facing, they provide direct access to internal networks, and the majority in this dataset were exploited before disclosure. For these assets, a comprehensive and continuously updated asset inventory is the foundation — organizations cannot protect what they cannot see. From that foundation, compensating controls and mitigation

layers become essential to reduce exposure during the window between exploitation and patch availability.

Infrastructure and backbone systems — middleware, application servers, backup infrastructure, network equipment — represent the deepest remediation challenge. Median close times for infrastructure vulnerabilities routinely extend into months: Cisco IOS XE at 232 days, F5 AuthBypass at 231 days, WebLogic RCE at 125 days, Veeam Creds at 116 days. Compare this to endpoint medians that consistently fall under 14 days. The gap is structural, driven by change windows, downtime constraints, and manual update processes. This is where the long tail lives. This is where the Manual Tax is heaviest.

Reporting a single blended remediation metric across all three categories can obscure meaningful differences in risk posture. Understanding where your organization falls on each axis — endpoint speed, perimeter visibility, infrastructure depth — is the first step toward a remediation strategy that addresses the actual exposure rather than the average.





Bridging the Gap: Mitigation as a Remediation Layer

For the infrastructure class where patch cycles are measured in weeks or months, patching alone cannot close the exposure window. The patch remains the definitive fix, but the timeline for these assets demands that teams deploy additional response layers — network-level containment, virtual patching,

host isolation, traffic blocking, and compensating controls — to reduce exposure while the patch works its way through change approvals and maintenance windows. When AWE is measured in hours and the patch cycle in months, the space between them must be filled with active mitigation.

The Wasted Head Start

Zero-day losses are, in a sense, technology failures — the patch did not exist. But the data reveals a costlier and more preventable failure mode: n-day vulnerabilities where defenders had substantial warning and still lost.

WebLogic RCE (CVE-2023-21839) was first exploited 71 days after disclosure. The average organization took 242 days to remediate. That two-month head start was not wasted by a hypothetical adversary — it was consumed by real ones. China-based cryptomining threat actor Water Sigbin (also tracked as the 8220 Gang) leveraged the vulnerability to deploy PowerShell-based loaders delivering XMRig cryptocurrency mining malware at scale. The same flaw was exploited by Prophet Spider, a financially motivated actor operating within the Russian ransomware-as-a-service ecosystem. Two distinct threat actors, spanning nation-state and criminal motivations, all exploiting a vulnerability for which a patch had been available for over two months.


UnRAR Directory Traversal (CVE-2022-30333) saw first exploitation at day 70; defenders averaged 178 days to close. The vulnerability attracted ShadowSyndicate — a financially motivated threat actor that emerged in mid-2022 with operational ties to multiple ransomware-as-a-service groups including CI0p, BlackCat/ALPHV, and Ryuk — as well as MalasLocker, a ransomware group that leveraged the flaw as part of its intrusion chain. Once again,

these were not sophisticated zero-day campaigns requiring nation-state resources. They were opportunistic actors exploiting known, patched vulnerabilities against organizations whose remediation processes had not delivered the fix.

These wasted head starts represent the highest-ROI targets for operational improvement — because the patch existed, the threat actors were known, the exploitation patterns were documented, and the only thing missing was the mechanism to apply the fix at speed. This is not an intelligence failure. It is an operationalization failure.

From Observation to Elimination

The data in this section proves that observing vulnerabilities is not the same as eliminating them. A Risk Operations Center (ROC) approach bridges that gap — shifting from passive monitoring of CVE counts to operationalized risk elimination: unifying asset context, quantifying exposure in business terms, and orchestrating remediation and mitigation at machine speed. When the dynamics of the attacker-defender asymmetry guarantee that manual processes will lose the majority of these engagements, the only viable response is to automate the entire cycle from detection through action.




4

Risk Mass — From Counting Vulnerabilities to Measuring Exposure

The metrics most organizations rely on to measure remediation performance share a common flaw: they count events. Traditional remediation metrics capture two things well — whether a vulnerability was closed, and how quickly. What they do not capture is the cumulative exposure an organization absorbed while the vulnerability remained open. A CVE patched in one day and a CVE patched in six months both register as “remediated” — but the organization that carried the latter lived with 180 times more



opportunity for exploitation. This is the equivalent of measuring flood damage by counting rainstorms rather than measuring the water level — the count tells you something happened, but nothing about the cumulative impact.

Severity scores account for how dangerous a vulnerability is. Risk Mass accounts for how long and how broadly that danger persisted.



We adopt the term Risk Mass — analogous to cumulative exposure metrics long established in epidemiology and reliability engineering — to bring a dimension into vulnerability management that CVE counts and MTTR both ignore: time. The formula is straightforward — the number of vulnerable assets multiplied by the number of days each remains exposed. The result is expressed in exposure-days, a single value that captures the total window an attacker had to exploit an environment. Mathematically, Risk Mass is the area under the remediation curve — a concept well understood in survival analysis, applied here to the specific problem of vulnerability exposure measurement. Operationally, it is the accumulated cost of every day your organization did not act.

A vulnerability affecting 400 assets and closed in one day produces 400 exposure-days. The same vulnerability left open for 100 days produces 40,000. CVE counts treat these as the same event. Risk Mass reveals them as fundamentally different risk postures.



Follina (CVE-2022-30190): Anatomy of 33,000 Exposure-Days

To demonstrate how Risk Mass works in practice, we applied it to one of the most consequential vulnerabilities of the past four years.

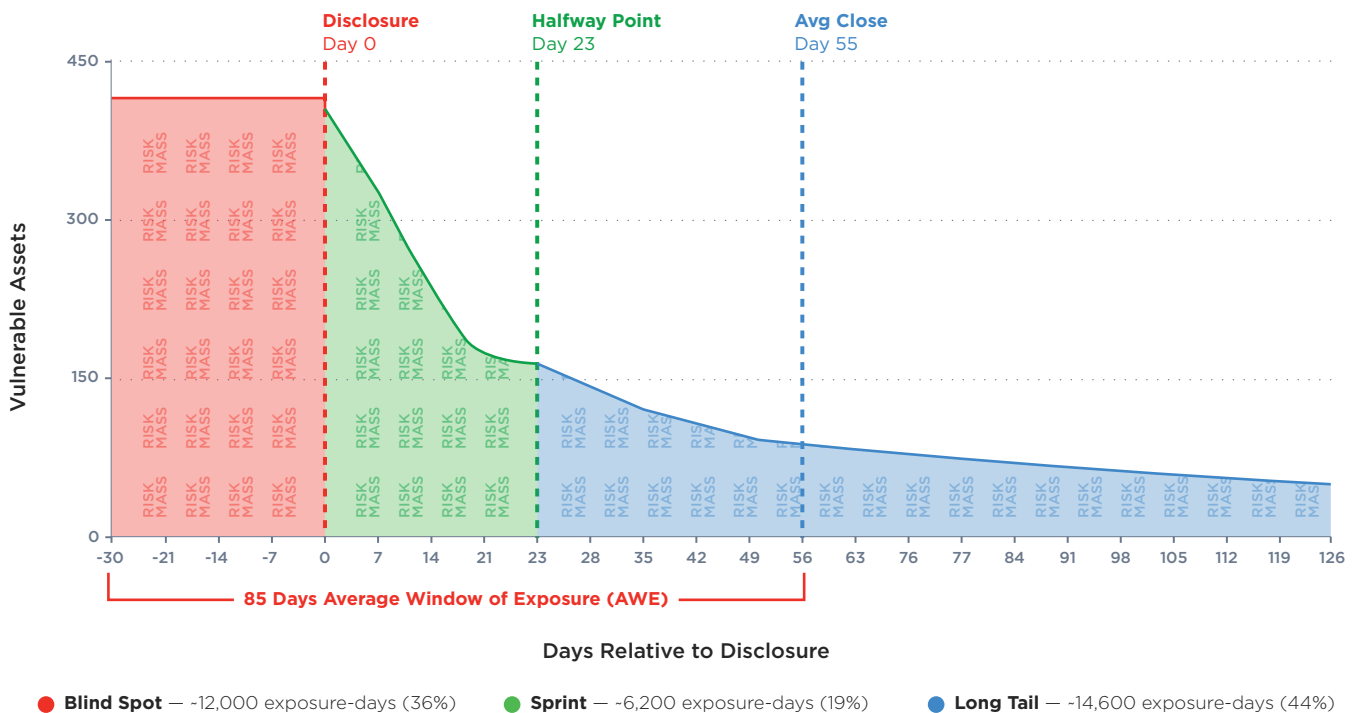
Follina is a remote code execution vulnerability in the Microsoft Support Diagnostic Tool that requires no macros and can be triggered through a specially crafted Word document. Follina's threat actor timeline mirrors its Risk Mass phases. During the Blind Spot, Chinese-affiliated

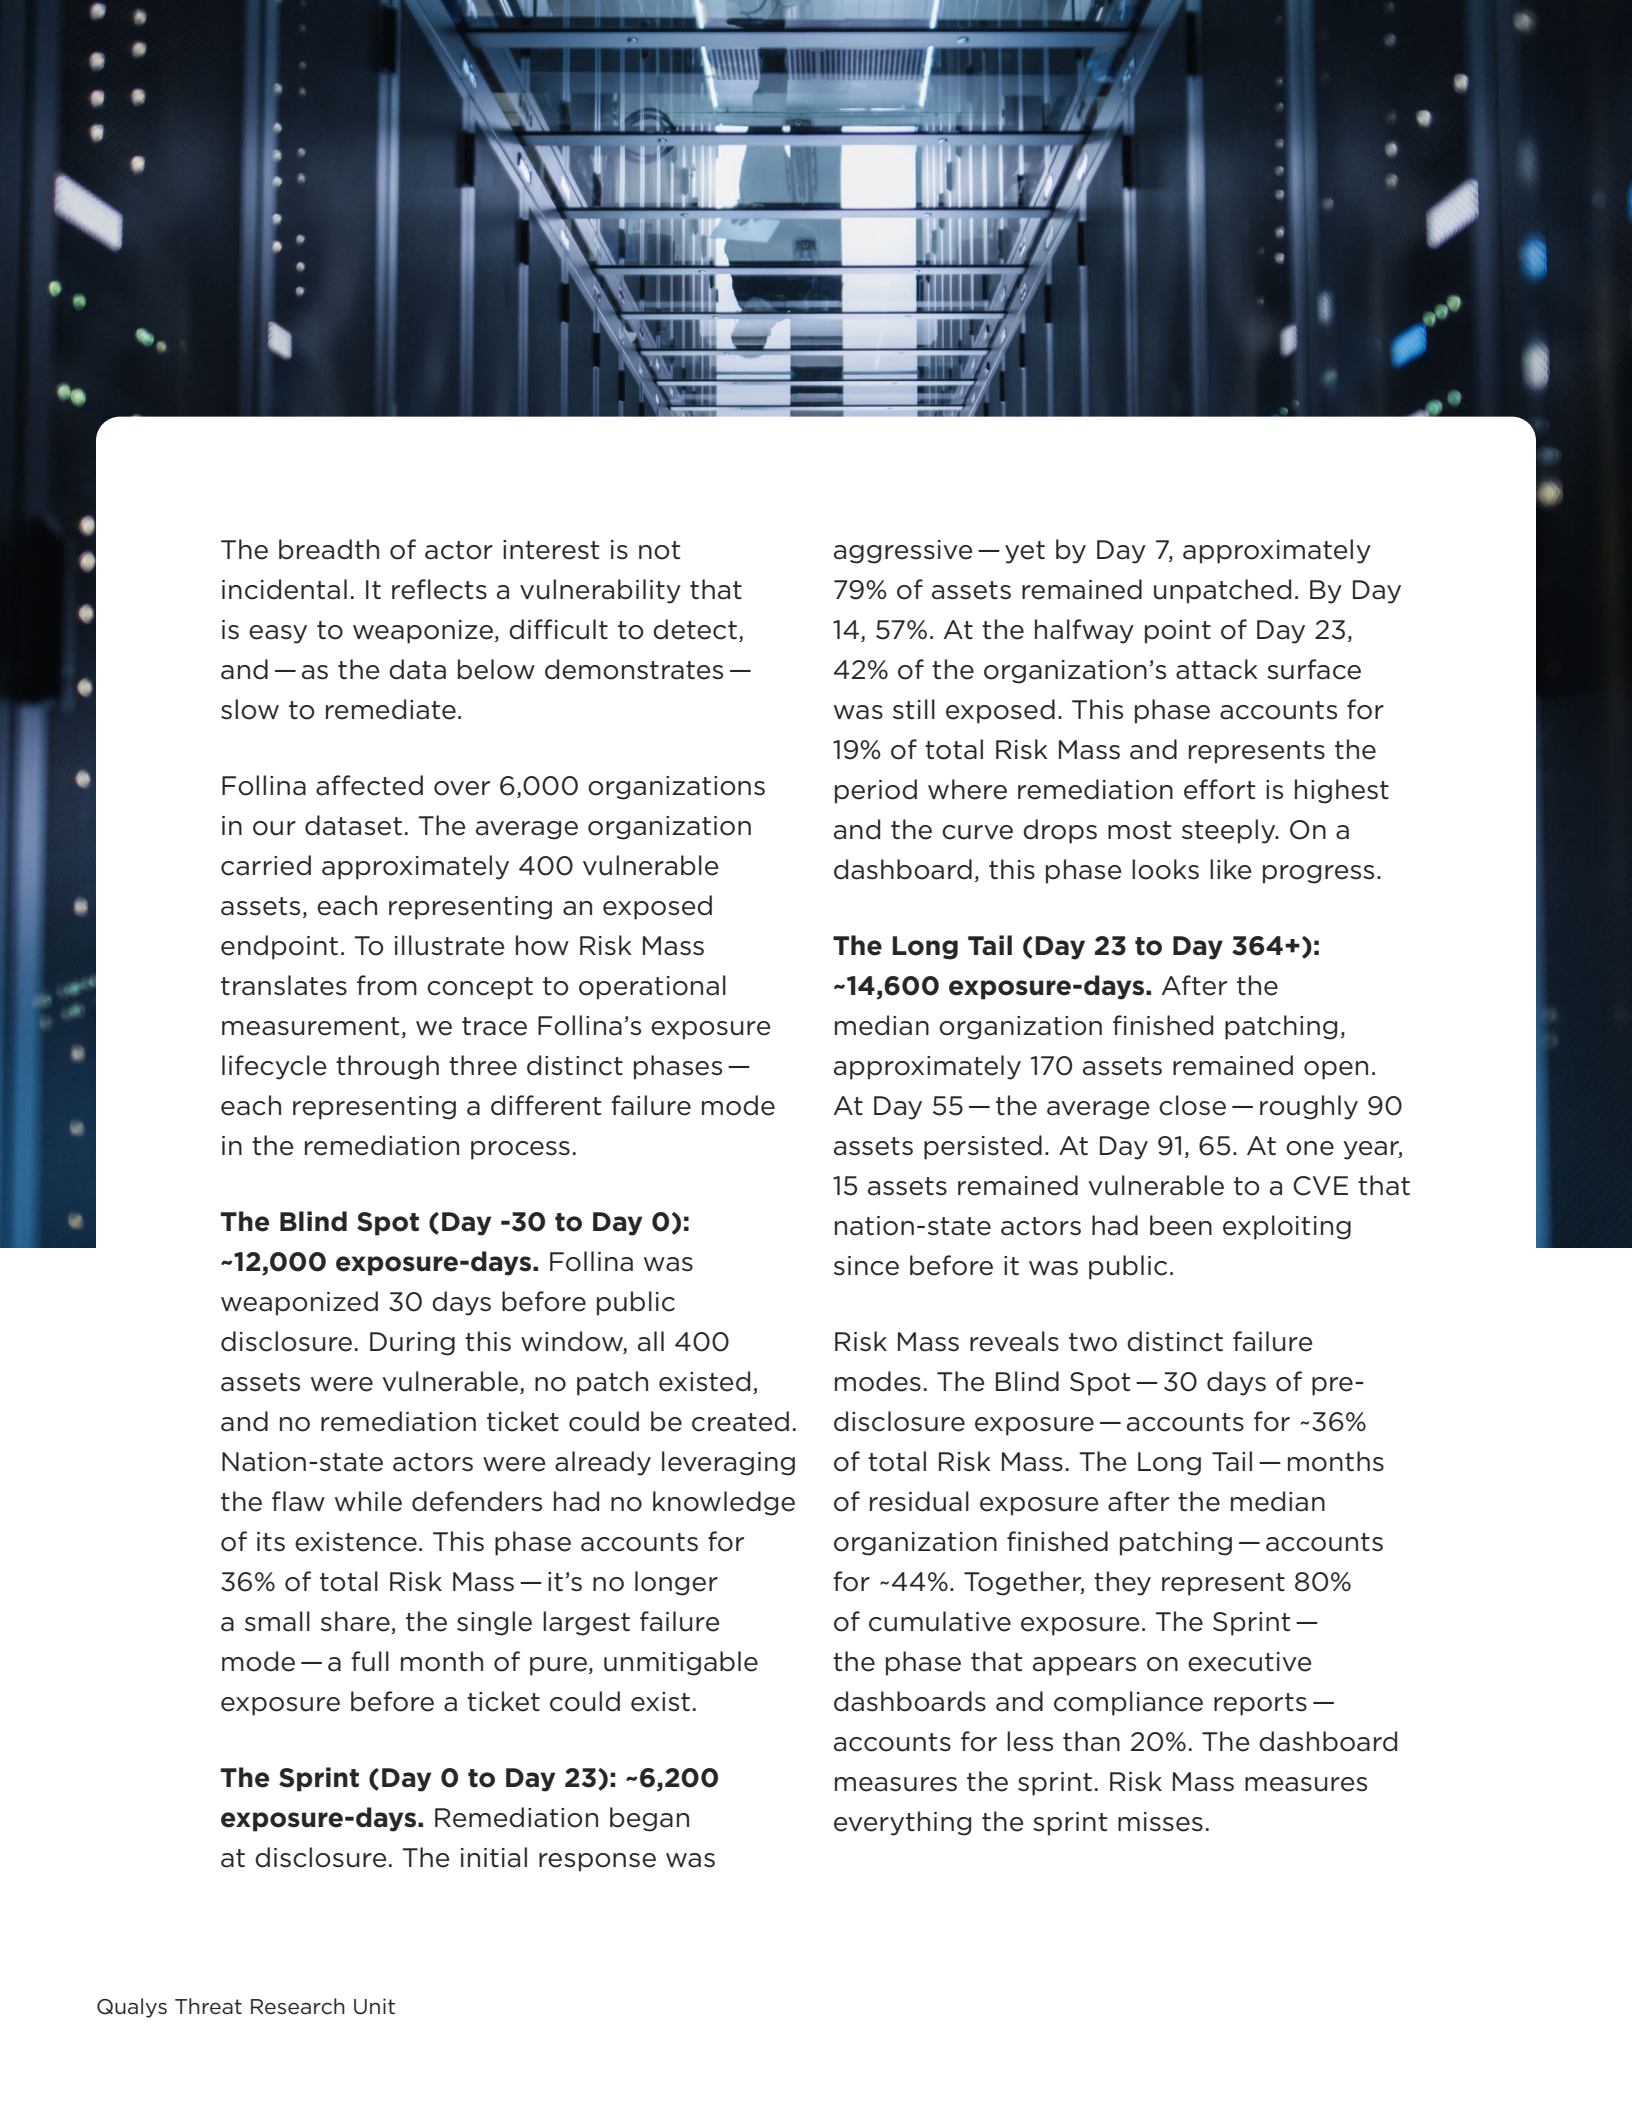
actors exploited the flaw using targeted lures, and TA413 deployed it against Tibetan organizations – weeks before disclosure. During the Sprint, Russian military intelligence (APT28, Sandworm) weaponized it against 500+ Ukrainian targets, while UAC-0098 delivered Cobalt Strike to Ukrainian infrastructure. In the Long Tail, commodity operators adopted it – Qakbot affiliates, AsyncRAT, Rozena – and ransomware groups (Clop, LockBit 3.0, BLOODY, Buhti) integrated it into active campaigns. Each phase attracted a different tier of adversary. The assets still open in the tail faced the widest and least predictable threat surface.

Follina (CVE-2022-30190) • Average Organization

Risk Mass: ~33,000 Exposure-Days from a Single CVE

~400 vulnerable assets x time exposed





The breadth of actor interest is not incidental. It reflects a vulnerability that is easy to weaponize, difficult to detect, and—as the data below demonstrates—slow to remediate.

Follina affected over 6,000 organizations in our dataset. The average organization carried approximately 400 vulnerable assets, each representing an exposed endpoint. To illustrate how Risk Mass translates from concept to operational measurement, we trace Follina's exposure lifecycle through three distinct phases—each representing a different failure mode in the remediation process.


The Blind Spot (Day -30 to Day 0): ~12,000 exposure-days. Follina was weaponized 30 days before public disclosure. During this window, all 400 assets were vulnerable, no patch existed, and no remediation ticket could be created. Nation-state actors were already leveraging the flaw while defenders had no knowledge of its existence. This phase accounts for 36% of total Risk Mass—it's no longer a small share, the single largest failure mode—a full month of pure, unmitigable exposure before a ticket could exist.

The Sprint (Day 0 to Day 23): ~6,200 exposure-days. Remediation began at disclosure. The initial response was

aggressive—yet by Day 7, approximately 79% of assets remained unpatched. By Day 14, 57%. At the halfway point of Day 23, 42% of the organization's attack surface was still exposed. This phase accounts for 19% of total Risk Mass and represents the period where remediation effort is highest and the curve drops most steeply. On a dashboard, this phase looks like progress.

The Long Tail (Day 23 to Day 364+): ~14,600 exposure-days. After the median organization finished patching, approximately 170 assets remained open. At Day 55—the average close—roughly 90 assets persisted. At Day 91, 65. At one year, 15 assets remained vulnerable to a CVE that nation-state actors had been exploiting since before it was public.

Risk Mass reveals two distinct failure modes. The Blind Spot—30 days of pre-disclosure exposure—accounts for ~36% of total Risk Mass. The Long Tail—months of residual exposure after the median organization finished patching—accounts for ~44%. Together, they represent 80% of cumulative exposure. The Sprint—the phase that appears on executive dashboards and compliance reports—accounts for less than 20%. The dashboard measures the sprint. Risk Mass measures everything the sprint misses.



5

The Filter — From Prioritization to Confirmation

The previous sections documented the speed at which adversaries exploit and the duration for which organizations remain exposed. This section addresses the next stage of the remediation pipeline: once organizations have successfully prioritized the vulnerabilities that matter, how can they add a final layer of deterministic confirmation to safely trigger autonomous remediation at machine speed?

The Prioritization Achievement

It is important to acknowledge what the industry has accomplished. Over the past decade, vulnerability management has evolved from indiscriminate patching based on CVSS severity alone to sophisticated, risk-informed prioritization that accounts for threat actor activity, asset context, and business criticality. Advanced risk scoring systems such as TruRisk have reached their highest adoption rates in history.

Organizations today have a clearer view of what matters than at any previous point in the discipline's history.

This is a genuine achievement — and it remains a necessary foundation. Nothing in this section argues against prioritization. What the data argues is that prioritization alone, even when executed well, leaves a critical gap between knowing what is dangerous and confirming what is actually exploitable in a specific environment. Closing that gap is the next evolution of the discipline, not a replacement of what came before.

The Sub-1% Reality

To understand why that gap matters, consider the raw mathematics of vulnerability disclosure.


In 2025, the industry disclosed 48,172 new vulnerabilities. Of those, our threat intelligence confirms that 357 — just 0.74% — were remotely exploitable, actively weaponized in the wild, and supported by working proof-of-concept code. This sub-1% reality has remained remarkably consistent even as total disclosure volume has nearly doubled:

This does not imply that the remaining 99% are irrelevant — compliance, defense-in-depth, and future-exploitation risk all warrant continued attention — but it defines the tier where speed of response is existential.

Threat intelligence and risk-based prioritization perform an essential function here: they reduce the operational universe from tens of thousands of CVEs to the fraction that represent real, weaponized risk. For organizations that have not yet adopted this layer, the difference between remediating 48,172 vulnerabilities and focusing on 357 is the difference between operational paralysis and a defensible program.

This pattern holds at an even broader scale. As of February 2026, the CISA Known Exploited Vulnerabilities (KEV) catalog — widely adopted across industries as a prioritization baseline and increasingly mandated as a compliance benchmark — contains 1,517 CVEs out of 315,354 published to date. That is 0.48% — roughly one in every two hundred vulnerabilities ever disclosed has been confirmed as actively exploited and cataloged by CISA. Whether through TruRisk scoring, CISA KEV alignment, or comparable frameworks, the conclusion is consistent: the operational universe that demands urgent action is a fraction of a percent of total disclosure volume. The challenge is no longer identifying that fraction. It is confirming which entries within it are actually exploitable in your environment — and acting on that confirmation before the window closes.

But even after this reduction — even when an organization has successfully narrowed its focus to the most dangerous sub-1% — a second filtration problem remains.



The Sub-1% Reality: In 2025, the industry disclosed **48,172** new vulnerabilities. Of those, our threat intelligence confirms that **357** — just **0.74%** — were remotely exploitable, actively weaponized in the wild, and supported by working proof-of-concept code. This **sub-1%** reality has remained remarkably consistent even as total disclosure volume has nearly doubled.


The Sub-1% Reality

Finding the Needle: 357 Weaponized Vulnerabilities in 48,172 in 2025

Each dot represents 1 CVE disclosed in 2025

● **Red dots** = remotely exploitable, actively exploited in the wild, with confirmed PoC.

357 OUT OF 48,172 DOTS = 0.74%



YEAR OVER YEAR

2022

25,043
131

disclosed
weaponized 0.52%

2023

28,816
185

disclosed
weaponized 0.64%

2024

39,964
218

disclosed
weaponized 0.55%

2025

48,172
357

disclosed
weaponized 0.74%

ALL-TIME

315,354
total CVEs

CISA KEV CATALOG

1,517
actively exploited (0.48%)

TAKEAWAY

**Prioritization reduces the haystack.
Confirmation finds which needles are real.**



The Confirmation Gap

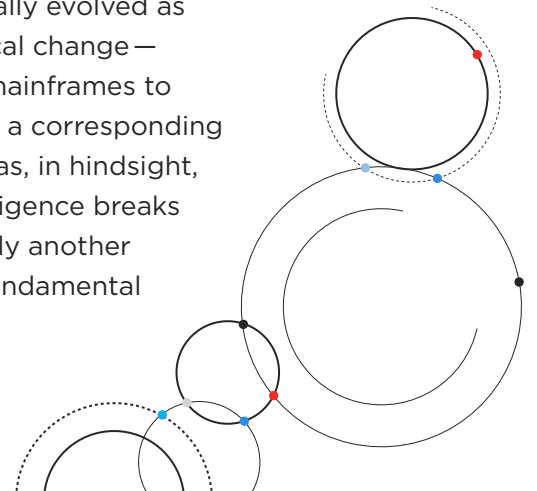
Risk-based prioritization — whether through TruRisk, CISA KEV alignment, or comparable frameworks — correctly identifies the vulnerabilities that represent the highest theoretical threat to an organization. This is exactly what it is designed to do, and the sub-1% filtering demonstrated above proves it works. The question that prioritization is not designed to answer is a different one: given this specific organization's deployed compensating controls, is this high-priority vulnerability actually exploitable right now?


That question can only be answered at the environment level. A vulnerability correctly scored as critical may be mitigated in practice by a Web Application Firewall blocking the exploit path, a service running in an inactive configuration, network segmentation isolating the asset from attacker-reachable infrastructure, or an endpoint detection rule that would intercept the payload on execution. These

controls do not reduce the vulnerability's severity — they reduce its exploitability in context. The vulnerability is real. The prioritization is correct. But the risk to this specific environment is lower than the score alone can express.

We refer to this residual gap as environment-adjusted exploitability — the difference between what is theoretically dangerous and what is confirmed exploitable given the controls in place. The final step before triggering automated remediation is confirming actual exploitability in context, so that machine-speed action is directed at validated exposure rather than theoretical risk.

Cybersecurity has historically evolved as a derivative of technological change — each new platform, from mainframes to mobile to cloud, produced a corresponding security imperative that was, in hindsight, predictable. Artificial intelligence breaks that pattern. It is not merely another surface to defend; it is a fundamental





restructuring of the threat landscape itself, one in which offensive actors are transitioning from human operators assisted by tooling to autonomous agents capable of identifying, weaponizing, and exploiting vulnerabilities at machine speed. The implications are severe: exploit timelines have compressed from weeks to hours, ushering in what practitioners now call -1 day — a condition in which weaponized exploits circulate before patches exist. Yet most enterprise security operations remain anchored to manual, human-in-the-loop workflows — analysts triaging thousands of theoretical vulnerabilities from scanners that cannot distinguish exploitable risk from statistical noise, teams mired in patching debates without evidence, and remediation cycles that move at human velocity against adversaries that increasingly do not.

This asymmetry defines the central challenge of the current era: until defenders can match the speed and autonomy of AI-driven offense, the gap will widen. It is within this context that

agentic AI orchestration emerges not as an incremental improvement but as an operational necessity. Agent Val, the agentic AI orchestration layer within Qualys Enterprise TruRisk Management, represents this shift — autonomously validating real exploitability in production through TruConfirm, confirming which attack paths are open and which are neutralized by existing controls, triggering targeted mitigation, and re-validating outcomes to deliver closed-loop proof of risk reduction. The imperative is no longer faster scanning or broader coverage; it is the elimination of human-speed bottlenecks from the defensive cycle entirely, replacing assumption-based prioritization with evidence-backed, machine-validated exposure management before the window between disclosure and exploitation closes for good.



From Probability to Higher Confidence


Closing the confirmation gap requires moving from version-based detection to exploit-based validation — replacing the question *“is this software vulnerable?”* with *“can this vulnerability actually be exploited in this environment, right now?”*

This is the function Qualys TruConfirm serves. Rather than matching software versions against advisory databases, TruConfirm operates as a safety-first, exploit-based confirmation engine that tests the actual execution path an attacker would use. It employs state-of-the-art, multimodal validation techniques — adjusted dynamically based on the nature and exploit characteristics of each vulnerability — to determine whether the full exploit chain succeeds against the target in its production configuration, without disrupting operations. Confirmed exploitable vulnerabilities are escalated with deterministic evidence — if the exploit chain executed, the finding is definitive.

The operational implication is direct. When combined with the multimodal detection described above — traditional version-based assessment to identify what exists, followed by TruConfirm to validate what is actually exploitable — organizations can reduce their confirmed remediation target list to a fraction of what passive scanning alone would indicate. The resources freed from ghost risk remediation can be redirected entirely toward the vulnerabilities that represent validated, proven exposure. It is important to note that deprioritization reflects confirmed current-state protection, not permanent dismissal — compensating controls reduce immediate exploitability but do not eliminate the underlying vulnerability, and changes in the control environment may re-expose previously mitigated risks.

The operational value of this layered approach becomes clear when viewed through the lens of the remediation physics documented in this report. TruRisk reduces the universe from 48,000 vulnerabilities to the sub-1% that represent real, weaponized risk — without this layer, operationalized remediation has no starting point. TruConfirm then provides the deterministic proof required for that final critical subset to safely trigger action. Prioritization tells you what deserves attention. Confirmation tells you what deserves action. Together, they form the complete pipeline that operationalized remediation requires.





6

Operationalize or Fail

The evidence in this report — 6.5x volume growth, a TTE of -7 days, and a remediation long tail that absorbs the majority of cumulative risk — points to a single conclusion: the problem is not speed. It is the operational model itself.

The Risk Operations Center

The response to these findings is not incremental. It is architectural. The scan-and-report model — where vulnerabilities are discovered, scored, ticketed, and manually routed through remediation


queues — was designed for an era of lower volume, longer exploit timelines, and acceptable human latency. That era has ended. What replaces it is what we define as a Risk Operations Center (ROC): a repeatable, end-to-end operational pipeline that executes the full remediation lifecycle at machine speed.

A ROC is not a dashboard, a team name, or a governance layer. It is an operational architecture built on three key capabilities that must function as a continuous, automated chain. Qualys Enterprise TruRisk Management operationalizes all three within a single platform.

The first is embedded intelligence. The intelligence imperative described in Section 1 reaches its operational expression here. In a ROC, threat intelligence does not arrive as a deliverable to be consumed—it arrives as machine-readable decision logic that feeds directly into the remediation pipeline. When a vulnerability is disclosed, the system does not wait for an analyst to assess it. It correlates the CVE against the organization’s asset inventory, evaluates whether affected assets are internal or externally facing, cross-references threat actor activity and relevance to the organization’s industry vertical, checks for known exploitation in the wild, assesses existing compensating controls, and weighs historical remediation precedent for similar vulnerability classes—determining whether action is required, and at what priority, in seconds rather than days.

The second is active confirmation. This is the noise filter that eliminates the ghost risk problem. A ROC applies multimodal verification—traditional methods such as version-based detection and configuration assessment to identify what exists, followed by active confirmation through TruConfirm, a Safety-First exploit-based validation engine, to verify what is actually exploitable given the organization’s deployed controls. The result is a dramatically compressed target list. Instead of remediating thousands of theoretically critical vulnerabilities, teams and automated systems focus on the fraction that represent confirmed, validated exposure.

The third is autonomous action. For confirmed risks, the remediation response must compress to a timescale that matches the threat—but the degree of automation should reflect the organization’s operational maturity and risk tolerance. On endpoints where automated patching infrastructure is established, organizations can move toward policy-driven deployment informed by historical success data—leveraging patterns from previous patch cycles, rollback rates, and environment-specific outcomes to build confidence in automated decisions over time. For infrastructure assets where patch cycles are constrained by change windows and downtime requirements, the system can recommend and, where authorized, deploy compensating controls—virtual patches, network-level containment rules, host isolation—that reduce exposure while the





definitive fix works through approval processes. And where patching cannot occur within the target window, these compensating controls are not interim measures — they become the primary line of defense, providing continuous mitigation until the definitive fix can be applied. The goal is not to remove human oversight entirely. It is to shift humans from executing every action to governing the policies that determine when and how actions execute — a model that scales with volume in a way that manual triage cannot.


The End of Risk Whack-a-Mole

The instinct to treat each new critical CVE as an isolated emergency — to rally the team, escalate the ticket, and sprint toward remediation — is understandable. It is also unsustainable. When vulnerability volume grows at 6.5x, exploit timelines collapse below zero, and the long tail of manual remediation absorbs two-thirds of cumulative risk, the reactive model does not scale. Every sprint becomes a triage. Every triage leaves a longer tail. The tail compounds.

The path forward is not to respond faster to individual vulnerabilities. It is to build a repeatable operational process — from intelligence ingestion through confirmation through automated action — that executes consistently, at scale, without requiring human intervention at every stage. The organizations in our dataset that are already winning the physics gap are not winning because they have larger security teams. They are winning because they have operationalized the end-to-end remediation lifecycle in a way that removes human latency from the critical path.

The dynamics of exploitation will continue to accelerate. Time-to-Exploit will not return to positive numbers. Vulnerability volume will not plateau. The only variable still within the defender's control is the speed and consistency of the remediation response — and the data in this report demonstrates, across a billion records and four years of evidence, that the only way to compress that response to the timescale the threat environment now demands is to operationalize it completely.





7

Appendix: Survival of the Vulnerable — Extended Analysis from the DBIR 2026 Partnership

This appendix is the result of a research partnership between the Qualys Threat Research Unit (TRU) and the Verizon DBIR team. [The Verizon 2026 Data Breach Investigations Report \(DBIR\)](#) includes a condensed view of our CISA KEV survival analysis, drawn from more than 1 billion anonymized vulnerability remediation records. The DBIR team encouraged us

to publish the extended analysis — the deeper breakdowns, additional context, and operational interpretation that could not fit within the published DBIR report. This appendix is that deeper dive. We thank the DBIR team for the partnership and the opportunity to contribute.

Section 0 — Methodology and Dataset Context

How We Measured KEV Remediation Survival

This analysis keeps unresolved vulnerabilities visible. Conventional remediation metrics report only on closed tickets, systematically understating true exposure. The survival analysis approach used here tracks every vulnerability instance over time — open or closed — ensuring the long tail of unremediated risk is never hidden behind closure-rate statistics.

Dataset Scope

The analysis draws on **1.01 billion** anonymized detection and remediation records spanning four consecutive DBIR reporting cycles (each running November 1 through October 31). All findings are restricted to CVEs listed in the CISA Known Exploited Vulnerabilities (KEV) catalog. Throughout this appendix, “Year 1” through “Year 4” refer to these four cycles from 2021 through 2025.

The underlying dataset is the same remediation telemetry analyzed in the above report, which uses calendar-year windows (January through December). Because this appendix aligns to the DBIR reporting convention (November through October), individual figures may differ slightly between the main report and this appendix. The directional findings and conclusions are consistent across both windowing approaches.

Data is measured at the level of **individual vulnerability instances** — a specific CVE

on a specific asset — capturing operational exposure at scale, not simply unique CVE counts. Records classified as ignored, disabled, deleted, or informational were excluded to focus on actionable findings. The dataset reflects a large, anonymized global sample; it is not a claim of total coverage across all organizations, sectors, or geographies.

Why Survival Analysis

The central question is not “how fast did resolved tickets close?” but rather: **“What percentage of known-exploited vulnerability instances remain open at each successive milestone?”**

Closure-time metrics — median or mean time to remediate — describe only vulnerabilities that were eventually fixed. They say nothing about those that were not. Survival analysis corrects this by tracking every instance from its starting point forward. At each weekly milestone, the curve reports the share of the original population still open. The result is a true exposure curve, not a speed-of-closure curve.

The Two Remediation Clocks

Remediation performance is measured against two independent clocks, each answering a distinct operational question.

The gap between these two curves reveals the delay between awareness and action — the interval during which exposure accumulates. Detection-relative timing reflects when the vulnerability appeared in scanner results, not necessarily when responsible owners reviewed or acted on it.

Clock	Starting Point	What It Measures
KEV-relative	Date CISA added the CVE to KEV	Response to an authoritative external exploitation signal
Detection-relative	Date the vulnerability was first detected in the environment	Operational execution speed after direct visibility

Proactive Remediation

A subset of instances is remediated before CISA adds the CVE to KEV. These pre-KEV closures are recorded as Day 0 remediations in the survival curves and reported as a standalone metric—a leading indicator of intelligence-driven remediation maturity. Year-over-year changes in this rate reveal whether organizations are improving at acting ahead of the formal exploitation signal.

Known Limitations

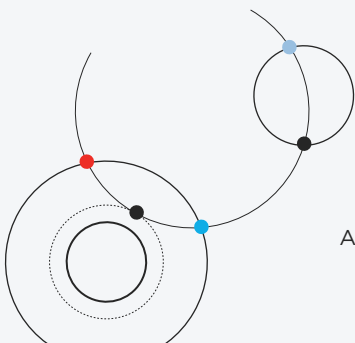
- **Year 1 is structurally different.** It coincides with the KEV catalog’s launch; organizations were encountering it for the first time. Year 1 findings are included for completeness but should be interpreted accordingly.
- **Late-window observation bias.** CVEs added near the end of a reporting cycle have shorter observation windows, which can make tail-end survival curves appear modestly better than reality.

Section 1— The KEV Survival Curve: The Remediation Treadmill Is Speeding Up

The four-year KEV survival curve is the single most important exhibit in this appendix. It shows a remediation system that improved materially from Year 1 through Year 3—then regressed in Year 4, even as the median detection-to-closure time held steady at 9 days for vulnerabilities that were eventually remediated. The regression did not happen because teams patched more slowly. It happened because volume grew faster than the system’s ability to keep up.

The Patching Capacity Paradox

Among instances that were eventually remediated, median detection-to-closure time held at 9 days in both Year 3 and Year 4. Organizations closed more instances in absolute terms than in any prior year. Yet at every survival milestone—Day 7, Day 28, Day 91, Day 182, Day 364—a larger share of the population remained open. The engine did not slow down. The conveyor belt sped up.



Metric	Year 1	Year 2	Year 3	Year 4
Total records	68.7M	120.8M	295.8M	527.3M
Closed	61.1M	115.9M	285.1M	480.2M
Still open at observation end	7.6M	4.9M	10.7M	47.1M
Closure rate	88.9%	96.0%	96.4%	91.1%

Values in the decline over 84 days column are percentage points.

Total records grew **7.7x** from Year 1 to Year 4 — an order-of-magnitude shift in workload within operating models not designed for it. The closure rate tells the same story from another angle: Year 2 and Year 3 closed over 96% of KEV instances; Year 4 fell to 91.1%. That five-point decline, at Year 4 scale, produced **47.1 million still-open instances** — a 4.4x increase over Year 3’s 10.7 million.


The Survival Curve: What Each Milestone Reveals

The KEV-relative survival curve tracks the percentage of instances remaining open at weekly intervals after CISA adds a CVE to the KEV catalog. This is not just a patching chart — it is an exposure-retention chart, showing whether organizations reduce attacker-relevant risk fast enough once the warning signal is public.

Year 3 was the strongest point in the four-year trend, with organizations closing KEV exposure faster at every milestone. Year 4 reversed that progress, moving back toward Year 2 levels. By Day 28, 35% of

instances remained open versus 27% in Year 3. The long tail hardened near 9% — roughly 47 million instances with no near-term path to resolution under the current operating model.

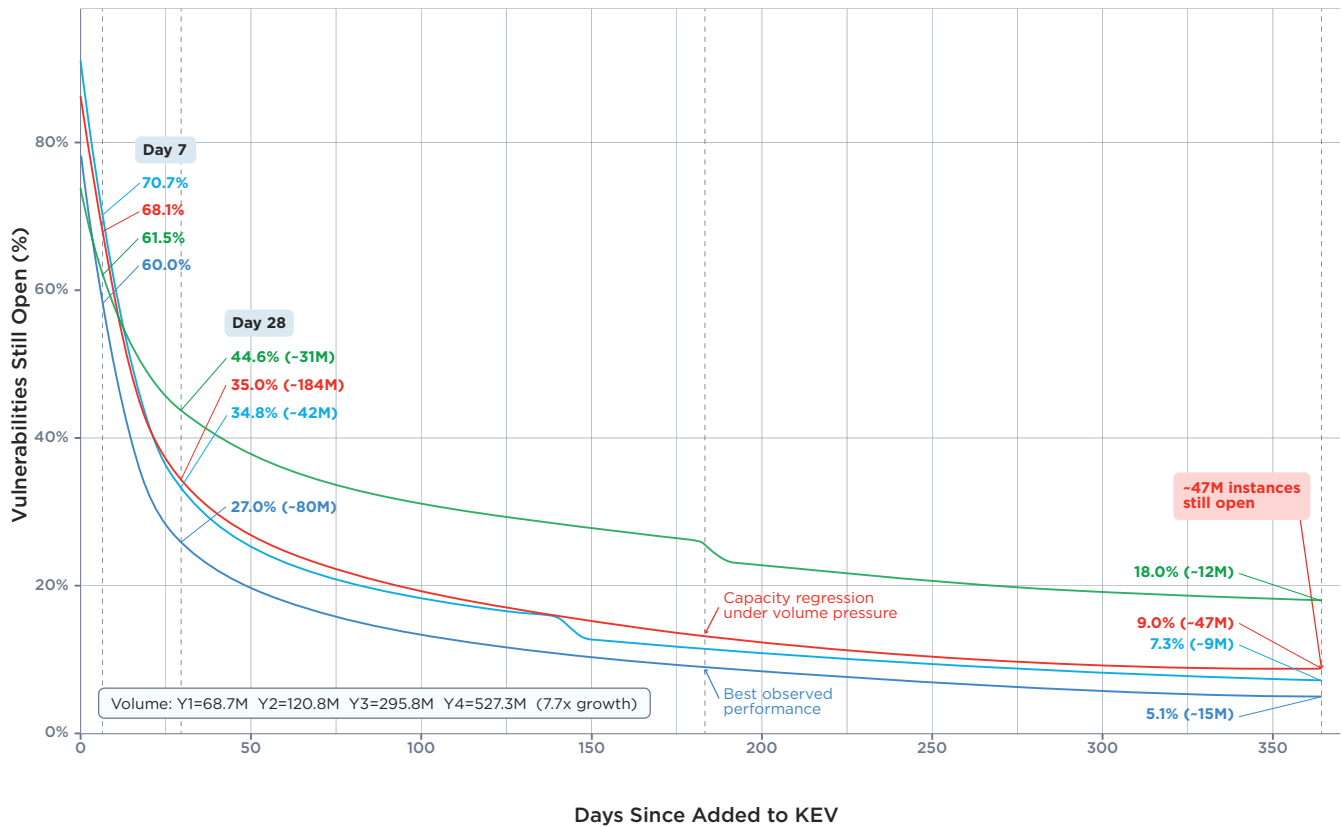
The practical reading: workload outgrew system capacity. The dataset grew nearly 8x while the number of distinct organizations grew modestly, meaning backlog expansion was driven primarily by higher exposure density per organization. The Day 28 backlog grew from approximately 31 million open instances in Year 1 to 184 million in Year 4. Teams improved their processes; volume overwhelmed the gains.



CISA KEV Vulnerability Survival Curve – 4 Year Comparison

The Remediation Treadmill is Speeding Up

2022 (Yr 1) 2024 (Yr 3)
2023 (Yr 2) 2025 (Yr 4)



The table below presents the five primary milestones across all four reporting cycles.

Milestone	Year 1	Year 2	Year 3	Year 4	Y3→Y4 Change
Day 7	61.5%	70.7%	60.0%	68.1%	+8.1 points
Day 28	44.6%	34.8%	27.0%	35.0%	+8.0 points
Day 91	32.4%	19.4%	14.6%	20.6%	+6.0 points
Day 182	25.8%	11.7%	9.2%	13.4%	+4.3 points
Day 364	18.0%	7.3%	5.1%	9.0%	+3.9 points

Values in the change column are percentage points.

Three patterns emerge from this data:

Pattern 1: Year 3 was the high-water mark.

Year 3 delivered the strongest remediation performance at every milestone. By Day 28, only 27.0% remained open; by Day 364, the long tail compressed to 5.1%. This was the year the industry came closest to demonstrating that the remediation model could keep pace with the KEV workload.

Pattern 2: Year 4 regressed toward

Year 2 levels. The largest regression is concentrated in the Day 7 through Day 91 window — the heart of the remediation cycle. The 8.1 percentage-point increase at Day 7 and 8.0 at Day 28 represent a real, system-wide shift in the curve, despite three additional years of tooling investment and process improvements.

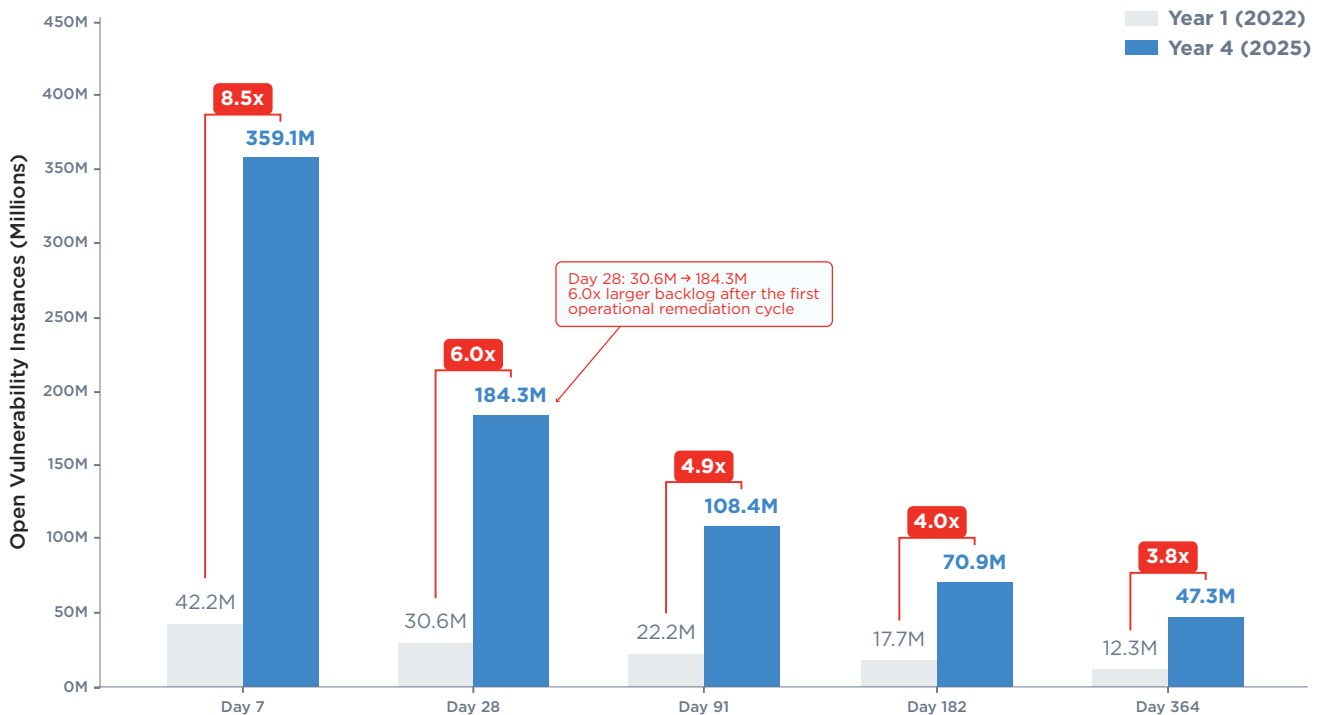
Pattern 3: The regression is widest where operational constraints bind hardest.

The Y3-to-Y4 gap narrows progressively from 8.1 points at Day 7 to 3.9 points at Day 364. Early milestones depend on rapid triage, immediate patching authority, and deployment capacity — all constrained by human decision-making. The fact that the regression is widest at these pressure points is consistent with a manual capacity constraint. The two-clock analysis later in this section separates that constraint into its components: when organizations became aware, and when they converted awareness into closure.

The Absolute Exposure Picture

Survival percentages are essential for cross-year comparison, but they can obscure the actual scale. The absolute open counts reveal the magnitude in terms that a risk leader can act on.

Percentages Improved, but the Open Backlog Multiplied



Milestone	Year 1	Year 2	Year 3	Year 4	Y3→Y4 Growth
Day 0	52.0M	110.7M	237.9M	452.4M	8.7x
Day 7	42.2M	85.4M	177.5M	359.1M	8.5x
Day 28	30.6M	42.0M	79.8M	184.3M	6.0x
Day 91	22.2M	23.4M	43.0M	108.4M	4.9x
Day 182	17.7M	14.1M	27.1M	70.9M	4.0x
Day 364	12.3M	8.8M	15.1M	47.3M	3.8x

The **Day 28** milestone is the most operationally significant checkpoint—approximating the window within which most organizations expect to have triaged, authorized, tested, and deployed a patch for a known exploited vulnerability. At this milestone, the absolute open backlog grew from 30.6M in Year 1 to 184.3M in Year 4—a 6.0x increase. Even maintaining Year 3 closure rates, sheer volume growth would have produced a substantially larger residual backlog.

By Day 364, **47.3 million instances** remained open in Year 4. These are not newly discovered findings awaiting triage. They are known exploited vulnerabilities, cataloged by CISA, visible in the environment, and still unresolved after twelve months. At this scale, the long tail is no longer a remediation queue—it is a structural feature of the exposure landscape.

Note: Day 364 is a fixed-horizon survival milestone, not a literal full-year observation for every CVE. CVEs added late in the reporting window have shorter observation

time. This is applied consistently across years but should be considered when interpreting the far-right tail.


Remediation Speed Versus Patching Capacity

The finding most likely to be misread: the remediation engine itself did not decelerate.


Median detection-to-closure held at 9 days in both Year 3 and Year 4. The average improved further, dropping from 32.0 to 23.6 days—an **8.4-day**, or **26%**, reduction. For vulnerabilities that were eventually remediated, organizations moved faster in Year 4 than in any prior year.

This is what makes the capacity thesis precise. The survival curve regressed not because defenders slowed down, but because the volume arriving on the conveyor belt outpaced the rate at which the belt could clear it. The engine is running at the same RPM; the load increased 7.7x.


The Load Exploded; the Median Closure Speed Did Not



Remediation Timeline (Days to Fix)



Year 4 saw a 78% increase in total KEV-linked vulnerability instances, yet the median time-to-close barely moved. Organizations patched individual vulnerabilities at roughly the same pace — but the queue behind each patch grew faster than teams could clear it. Speed held; capacity did not.




Metric	Year 1	Year 2	Year 3	Year 4
Median from KEV addition (days)	13	15	10	13
Median from first detection (days)	11	12	9	9
Average from KEV addition (days)	98.4	56.7	37.6	36.5
Average from first detection (days)	79.5	48.7	32.0	23.6

The distinction between speed and capacity determines the correct strategic response. If the engine had slowed, the answer would be operational improvements – better tooling, tighter SLAs. But the engine did not slow. The answer is structural: the remediation model itself must change to accommodate the volume and velocity of the modern KEV workload.

The Long Tail: Structural Residue, Not Slow Progress

The behavior of the survival curve between Day 280 and Day 364 reveals whether the remaining open population is still being actively worked on or has reached a structural floor where remediation has effectively ceased.

The Year 4 Long Tail Nearly Stopped Moving



Year	Day 280	Day 364	Decline Over 84 Days
<i>Year 1</i>	19.8%	18.0%	1.8 percentage points
<i>Year 2</i>	8.7%	7.3%	1.4 percentage points
<i>Year 3</i>	6.4%	5.1%	1.3 percentage points
<i>Year 4</i>	9.7%	9.0%	0.7 percentage points

Year 4 exhibits the flattest long tail in the dataset. Over 84 days — nearly three months — only 0.7 percentage points moved from open to closed, representing approximately **47 million instances** that are no longer progressing through the remediation pipeline.

These are not vulnerabilities awaiting discovery or lacking a remediation signal. They are KEV-listed, detected, and visible within vulnerability management programs — yet open after a full year. The long tail represents structural residue: findings where system criticality, ownership ambiguity, legacy architecture, change-window restrictions, or insufficient deployment automation have made remediation impractical under the current operating model. When the curve stops declining, the system has declared — through behavior, if not policy — that these vulnerabilities will not be fixed through normal workflow.

Strategic Synthesis

The remediation system improved — genuinely, measurably — from Year 1 through Year 3: faster closure, compressed long tails, stronger performance at every milestone. Year 4 broke the trajectory. Not

because investments failed, but because the workload scaled past the capacity those investments created. Volume grew 7.7x. The Day 28 backlog grew 6.0x. The long tail hardened to its flattest shape — 0.7 percentage-point movement over 84 days across 47 million structurally stuck instances.

Speed alone is no longer sufficient. A process that performs well on closed records can still fall behind when total workload reaches 527 million records, accumulating exposure at every milestone that matters. The implication is architectural, not incremental: organizations that cannot shift from manual, ticket-driven remediation to an operationalized, machine-speed pipeline will see this gap widen with every reporting cycle. The main report's Risk Operations Center (ROC) framework addresses this directly.

The survival curve proves the central paradox: defenders are operating at the same speed, but the volume of known exploited vulnerabilities has exceeded the patching capacity of the current model. The answer is not a faster engine. The answer is a different machine.

Section 2— Proactive Defense at Scale: 63.7 Million and Still Not Enough

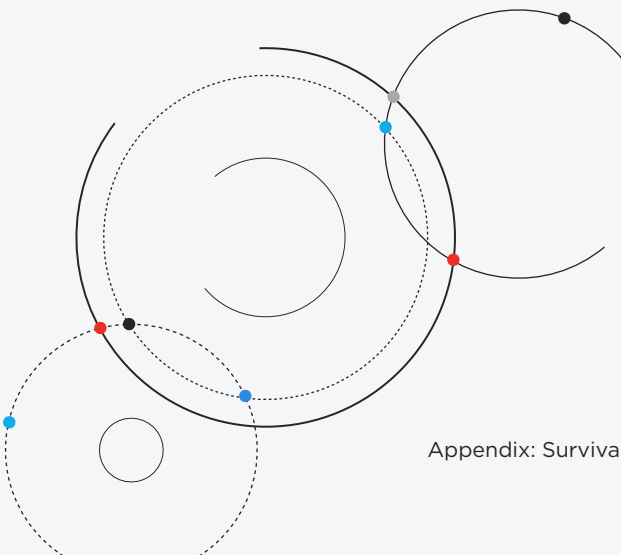
In Year 4, defenders proactively remediated **63.7 million** vulnerability instances before CISA cataloged them—a 30% increase from 48.9 million in Year 3. By any operational measure, proactive defense has never been stronger. Yet the proactive remediation rate fell from 16.6% to 12.1%. Read in isolation,

that suggests regression. Read against the absolute numbers, it reveals volume pressure: the total KEV workload grew 78%, nearly three times faster than the proactive count.

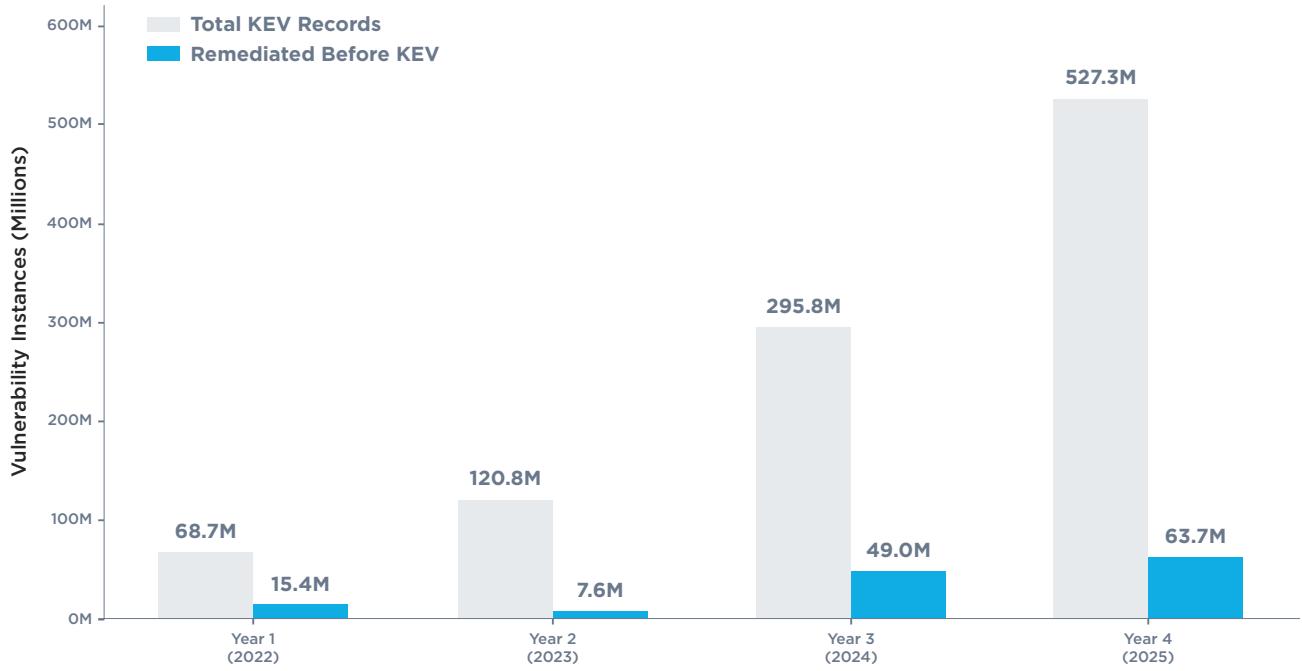
This is not a collapse of proactive behavior. Even the most mature defenders, operating with the best available intelligence and prioritization, cannot scale proactive action linearly against a workload that nearly doubles every reporting cycle.

The Proactive Remediation Record

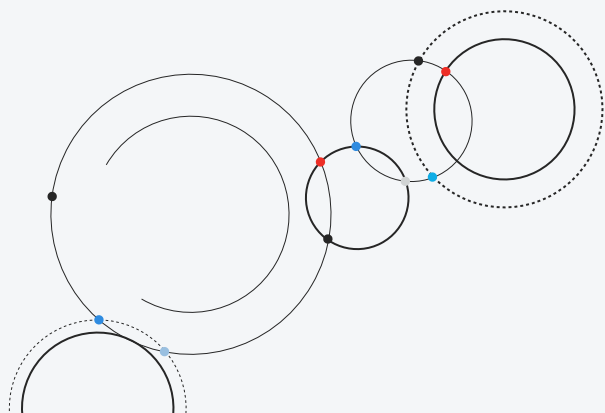

Metric	Year 1	Year 2	Year 3	Year 4
Total KEV Records	68.7M	120.8M	295.8M	527.3M
Remediated Before KEV Listing	15.4M	7.6M	49.0M	63.7M
Proactive Rate	22.5%	6.3%	16.6%	12.1%
KEV Workload Growth (YoY)	—	+75.8%	+144.9%	+78.3%



Proactive Remediation Grew, but the Share Fell



Proactive Remediation Rate




Year 1’s elevated rate (22.5%) reflects early catalog conditions: a smaller KEV list and a higher concentration of mature organizations in the initial dataset. Year 2 collapsed to 6.3% as the catalog expanded and the instrumented population broadened. Year 3 recovered to 16.6% as organizations matured their KEV ingestion workflows. Year 4 slipped to 12.1% under renewed volume pressure.

The Year 3-to-Year 4 transition isolates the volume effect precisely:

Metric	Year 3	Year 4	Change
Proactive Rate	16.6%	12.1%	-4.5 pp
Proactive Absolute Count	48.9M	63.7M	+30.2%
Total KEV Records	295.8M	527.3M	+78.3%

The value noted as “pp” in the change column denotes percentage points.

Proactive Defense Improved — But the Workload Outpaced It



The proactive count grew 30%. The total workload grew 78%. The rate declined because the workload expanded 2.6 times faster than proactive output.

What Proactive Defense Reveals About Organizational Maturity

The organizations driving proactive remediation have access to the same CVE disclosures and vendor advisories as every other defender. What distinguishes them is how they operationalize intelligence: leveraging risk-based prioritization, embedded threat context, and advanced scoring systems to act on exploitation signals before an authoritative catalog confirms what they already assessed to be true. Rather than waiting for CISA’s formal KEV addition — which may lag initial exploitation by days, weeks, or months — these organizations route likely-exploitable vulnerabilities into remediation workflows upstream of the signal.

The 63.7 million pre-KEV remediations in Year 4 are not accidental closures. They are the output of deliberate, intelligence-driven prioritization by organizations that have


moved beyond reactive patching. They are the proof that compressing the survival curve is possible. They are also, at present, the minority.

The Day 7 Ceiling: A Physical Constraint on Remediation Speed

The first-week window provides the most direct measurement of how much KEV workload the system can reach under maximum urgency.

On the detection-relative clock — which removes upstream visibility latency — the range tightens to 58–63% still open at Day 7 across all four years. At their operational best, organizations close 30–40% of KEV instances within the first week after detection. This ceiling has not moved materially despite three years of additional investment in tooling, process development, and regulatory pressure.

The First-Week Ceiling Did Not Move



Year	Day 7 Open (KEV Clock)	Day 7 Open (Detection Clock)
Year 1	61.5%	62.1%
Year 2	70.7%	63.5%
Year 3	60.0%	58.3%
Year 4	68.1%	60.2%

The stability of this range is itself a finding. It behaves less like a performance gap and more like a physical constraint—a “speed of light” for human-driven remediation. Change-control windows, testing requirements, deployment sequencing, and business-continuity dependencies impose a structural floor on how quickly patches move from “ready” to “deployed” across a heterogeneous enterprise. The implication: 60–70% of every new KEV cohort enters the backlog by default, regardless of urgency, staffing, or intent. The only way to change the proportion is to change the model—from manual, change-window-dependent remediation toward automated, policy-driven deployment that acts within hours of detection.

What the Data Proves

The organizations contributing to Year 4’s 12.1% proactive rate have demonstrated that acting ahead of the authoritative exploitation signal is possible—but it requires comprehensive detection coverage, embedded threat intelligence, risk-based scoring, and automated workflows that convert prioritization decisions into deployed patches without manual intervention at every stage.

The evidence shows mature programs are scaling proactive output: 63.7 million instances, up 30%. But linear scaling cannot outrun exponential workload growth. The rate fell because the workload grew faster. The Day 7 ceiling held because deployment remained manual. The long tail persisted because the constraints preventing remediation at Day 28 also prevent it at Day 364.

The path forward is transforming proactive defense from a maturity indicator into an operational default—converting intelligence into remediated assets at machine speed before exposure accumulates. The intelligence to identify the right vulnerabilities early exists. The remaining gap is execution—and, critically, confirmation. Prioritization identifies what deserves attention; environment-level validation confirms what is actually exploitable given deployed compensating controls, so that machine-speed action targets validated exposure rather than theoretical risk. Together, prioritization, confirmation, and autonomous remediation form the complete pipeline that operationalized remediation requires.




About Qualys

Qualys, Inc. (NASDAQ: QLYS) is a pioneer and leading provider of disruptive cloud-based security, compliance and IT solutions with more than 10,000 subscription customers worldwide, including a majority of the Forbes Global 100 and Fortune 100. Qualys helps organizations streamline and automate their security and compliance solutions onto a single platform for


greater agility, better business outcomes, and substantial cost savings. For more information, please visit qualys.com.

Qualys, Qualys VMDR® and the Qualys logo are proprietary trademarks of Qualys, Inc. All other products or names may be trademarks of their respective companies.





Qualys Threat Research Unit



www.qualys.com/tru

About the Qualys Threat Research Unit (TRU)

The Threat Research Unit (TRU) is the research arm of Qualys. The TRU team's focus is devoted to vulnerability, compliance, malware, and threat actor research with the goal of providing world-class security intelligence, detection data, and guidance for the Qualys Enterprise TruRisk™ Platform.


New technologies are revolutionizing lives and economies around the world.

Cyberthreats are growing at a similar pace, endangering access to the services that improve lives everywhere. By empowering customers and stakeholders for IT, security and compliance with TRU research and insights, together we can secure and defend the digital world from bad actors who create chaos and erode trust. We are the Qualys Threat Research Unit. Our shield is your shield.

Research Partner For

Verizon Data Breach Investigations Report





Qualys Threat
Research Unit

