Qualys Threat Research Unit

The Broken Physics of Remediation

The future of risk management is autonomous.

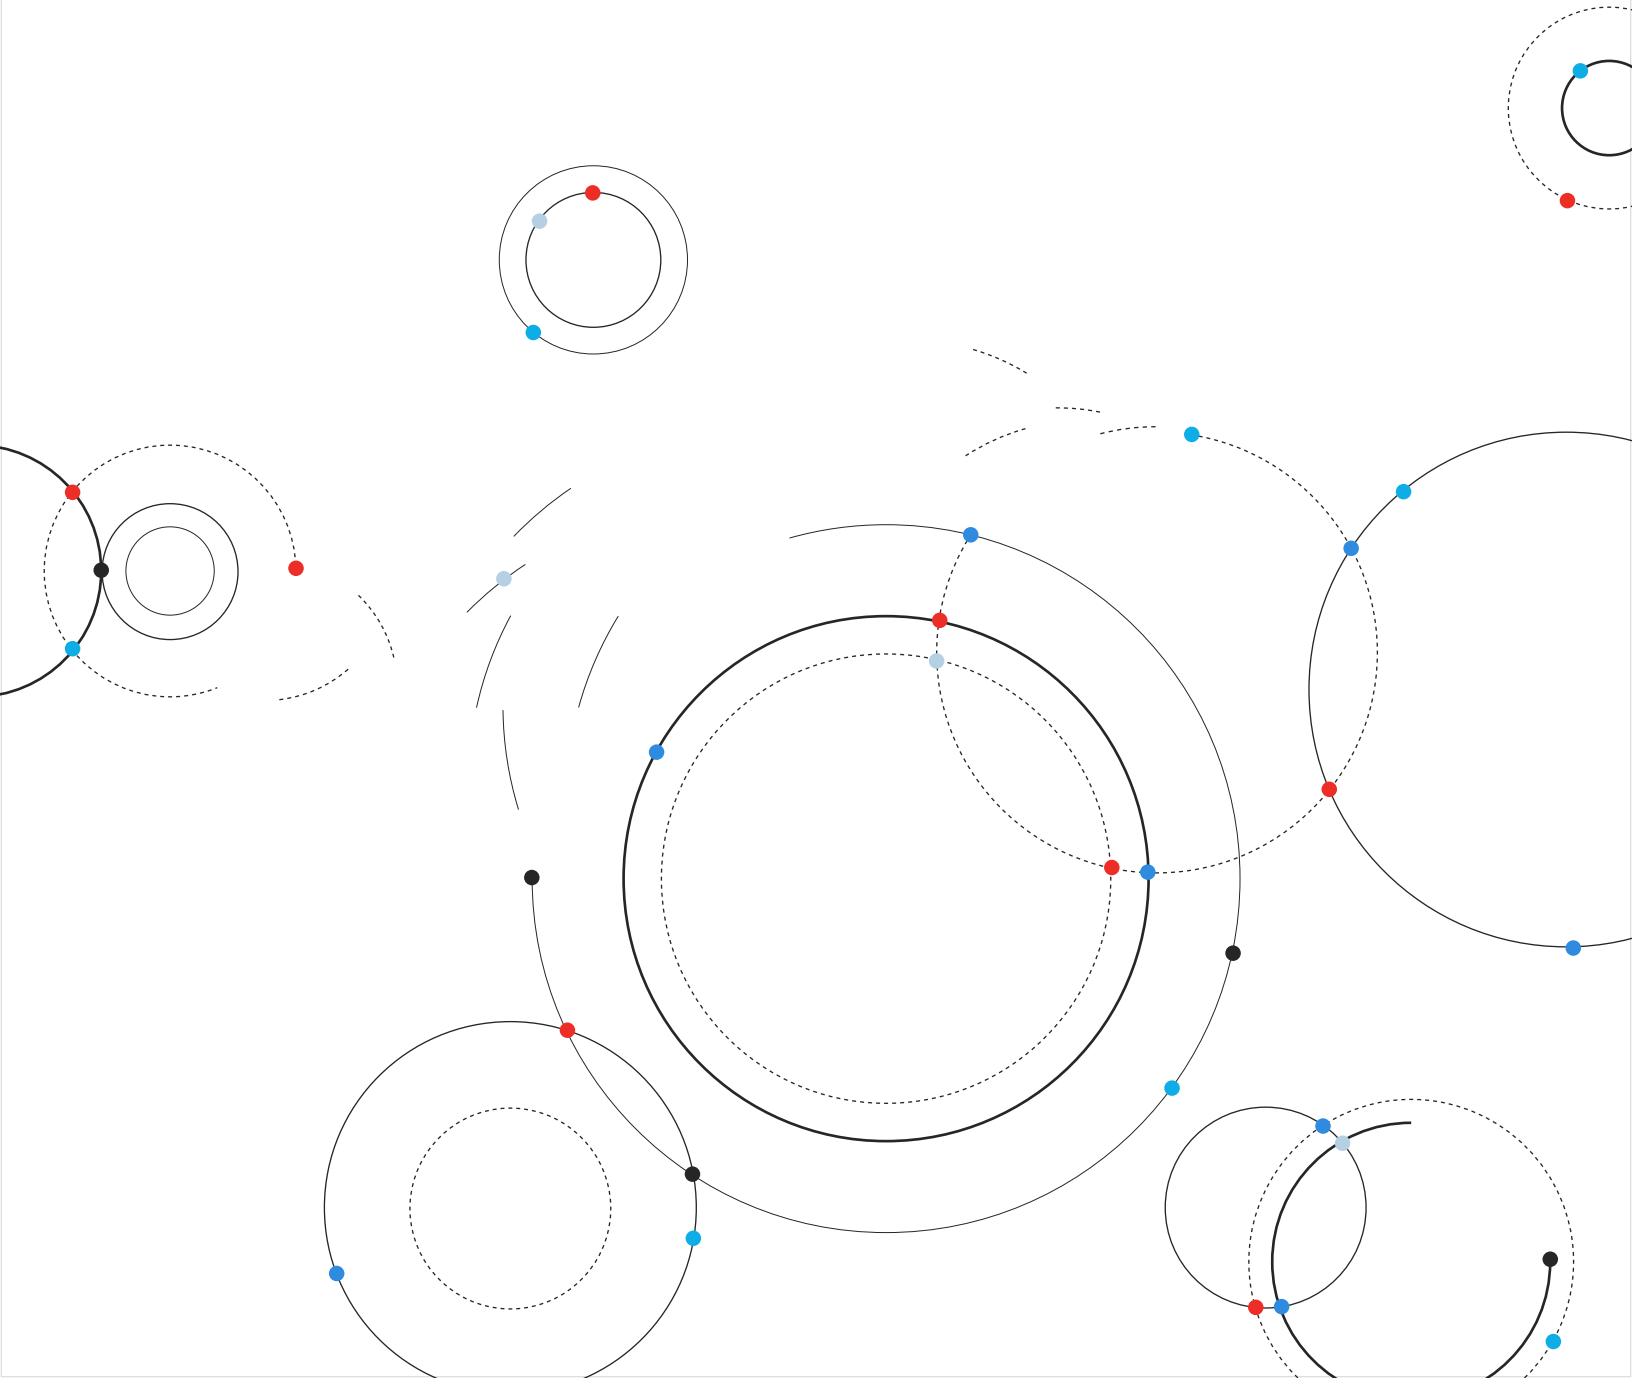

Foreword

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 one day 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.

Sumedh Thakar
President and CEO
Qualys

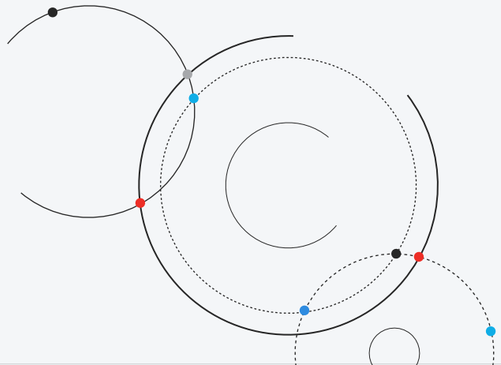

Introduction

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

Table of Contents

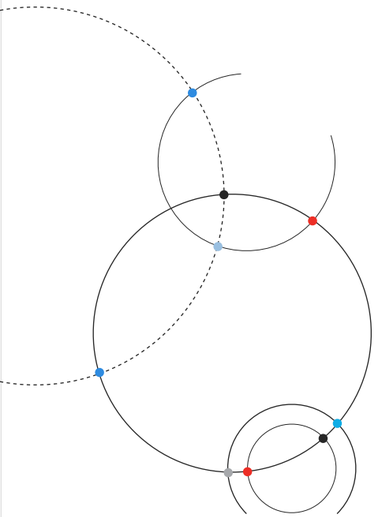

Qualys Threat Research Unit

Executive Summary

According to Google Threat Intelligence Group (GTIG) analysis of 2024 Time-to-Exploit Trends, published in September 2025, the average Time-to-Exploit (TTE) has collapsed to -1 Days. 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.

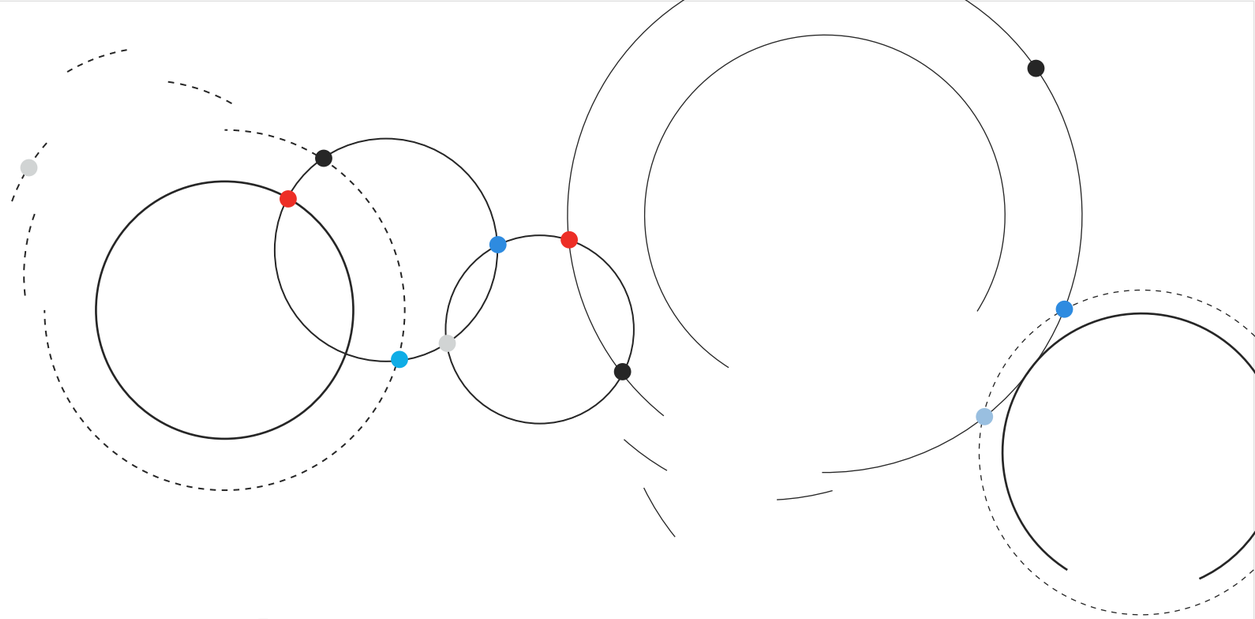

1 The End of Human-Scale Remediation

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.

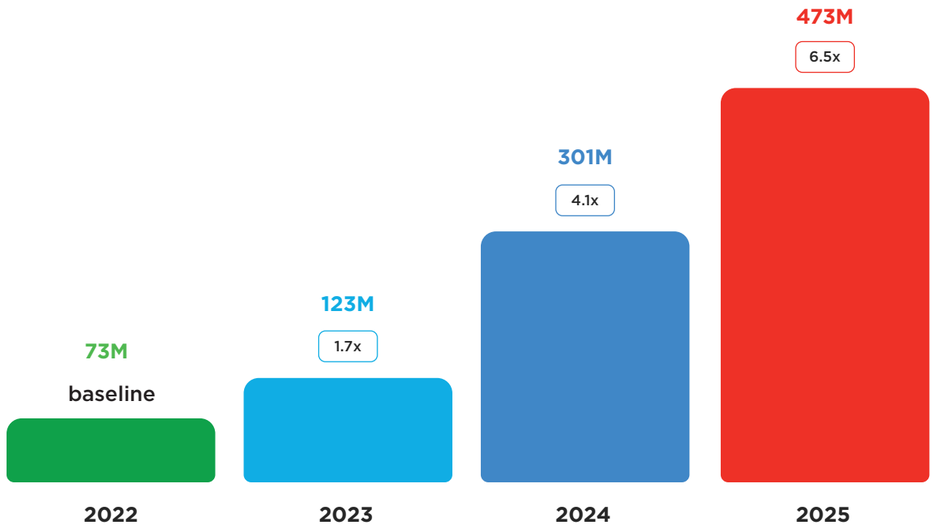

The volume of closed vulnerability events grew **6.5x** in four years — from approximately **73 million** in 2022 to **473 million** in 2025 — driven by both rising vulnerability volume and broader organizational coverage.

The 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.

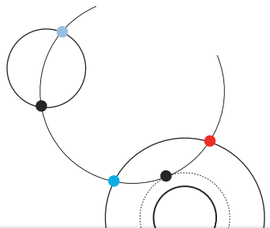

Even controlling for that growth, the percentage of critical CISA KEV vulnerabilities still open at Day 7 did not improve. It got worse — rising from 56% in 2022 to 63% in 2025.

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.

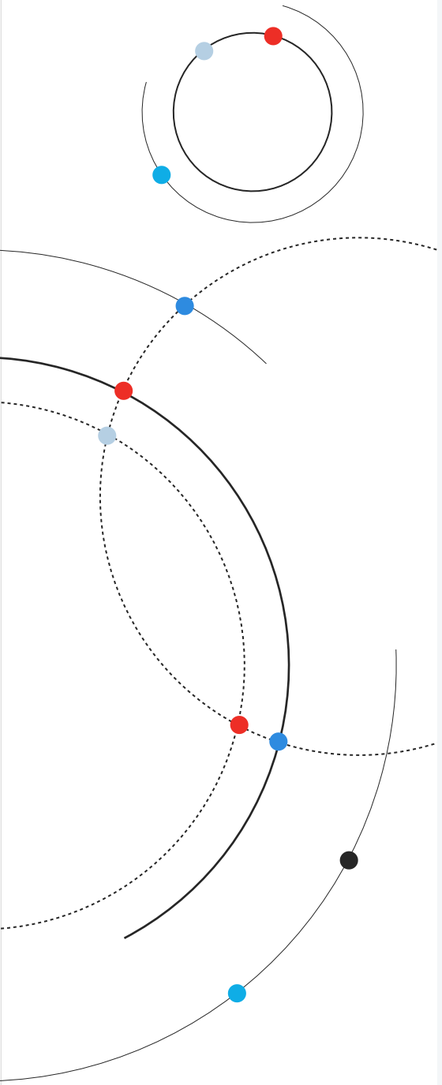

The Intelligence Imperative

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 -1 day, 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.

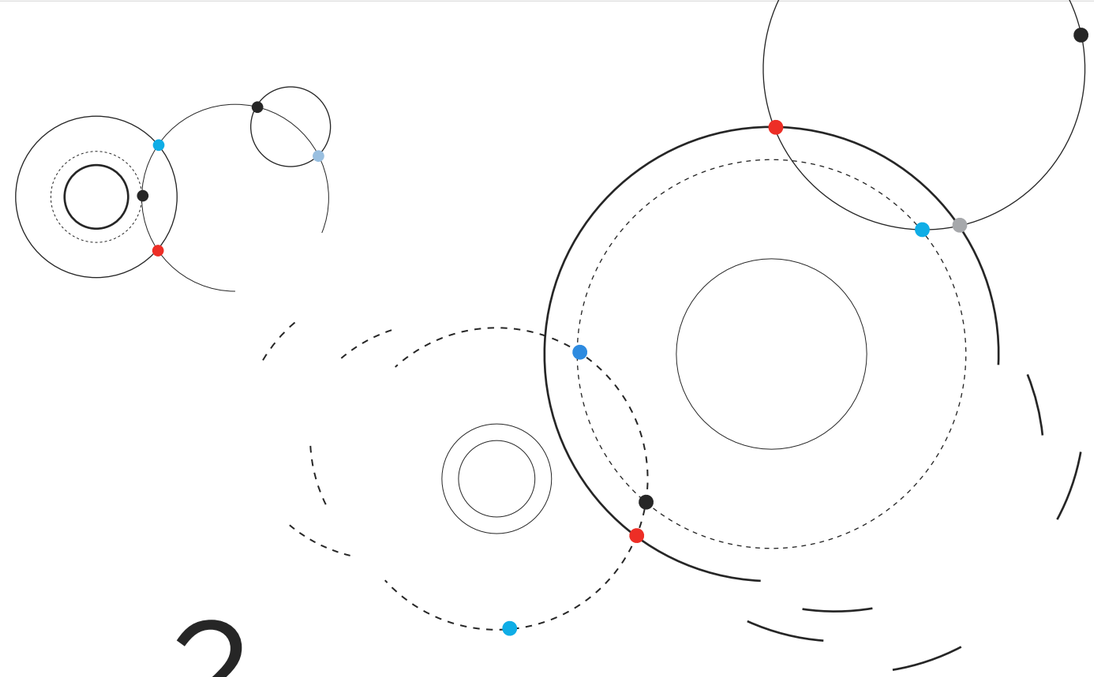

2 The Broken Physics of Remediation

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.

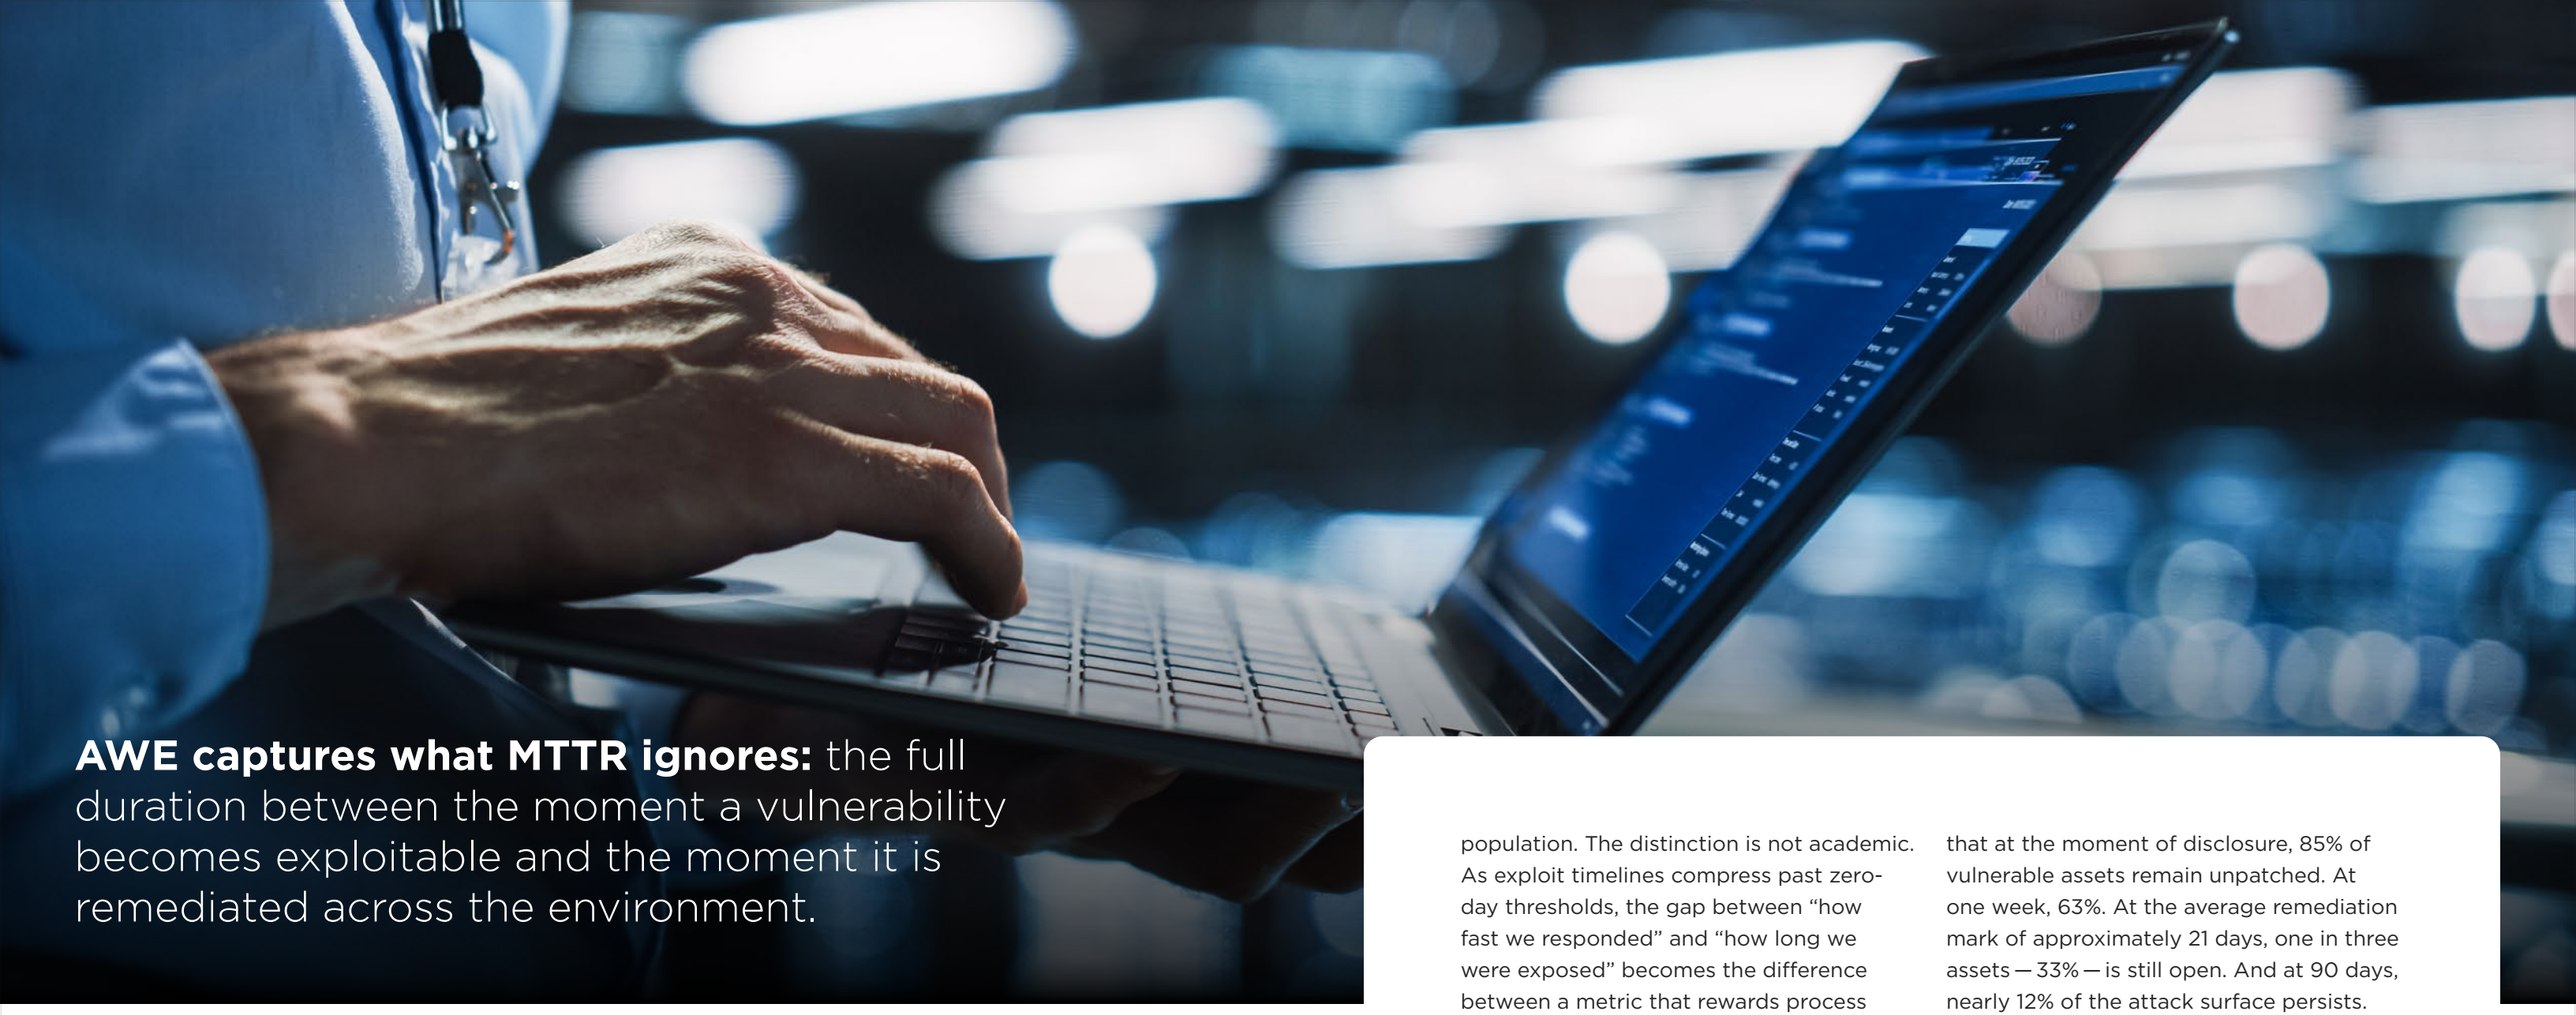

Average Window of Exposure

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.

GTIG's 2024 analysis of 112 exploited-in-the-wild vulnerabilities found that the average Time-to-Exploit has collapsed to -1 day — meaning that for vulnerabilities confirmed as actively exploited, weaponization is now occurring, 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: Average Window of Exposure (AWE)

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

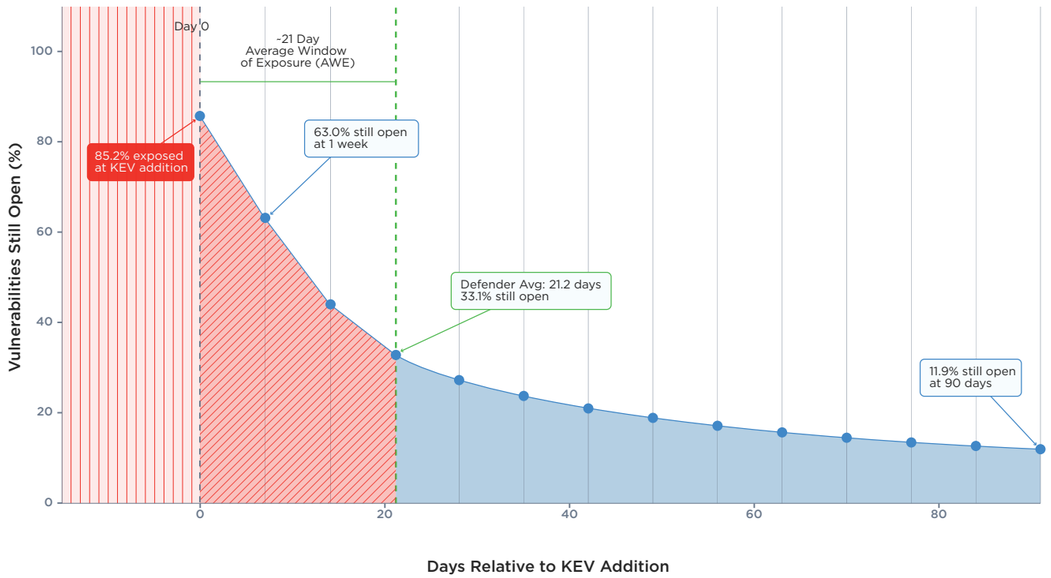

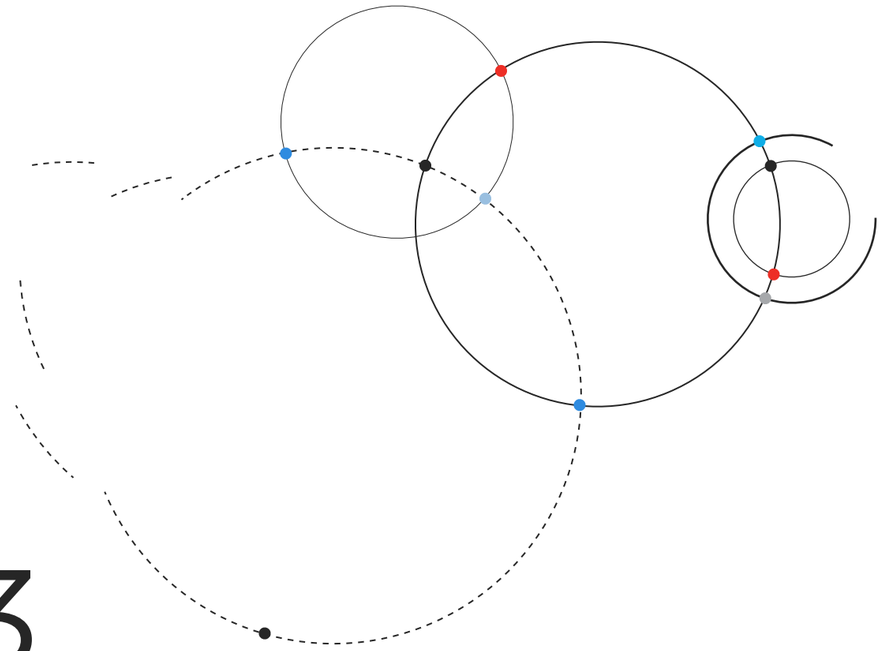

3 The Physics Gap — Attacker Speed vs. Defender Speed

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.

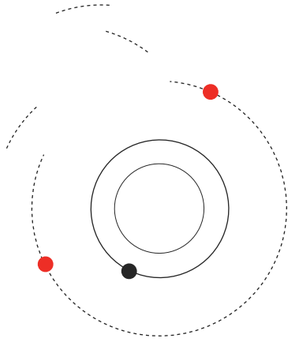

Reading the Physics Gap

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 Physics Gap: Attacker Speed vs. Defender Speed

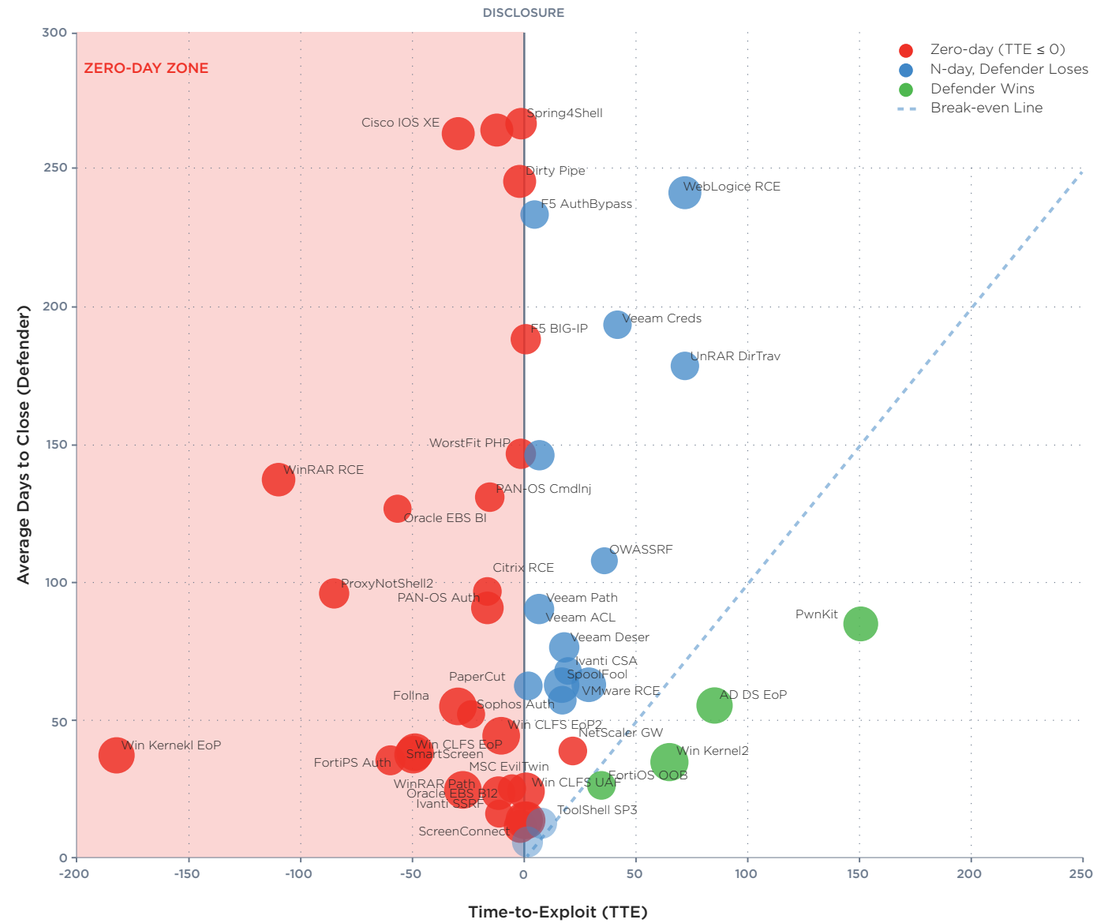

The Zero-Day Zone

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 RCE (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.

Qualys Threat Research Unit

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

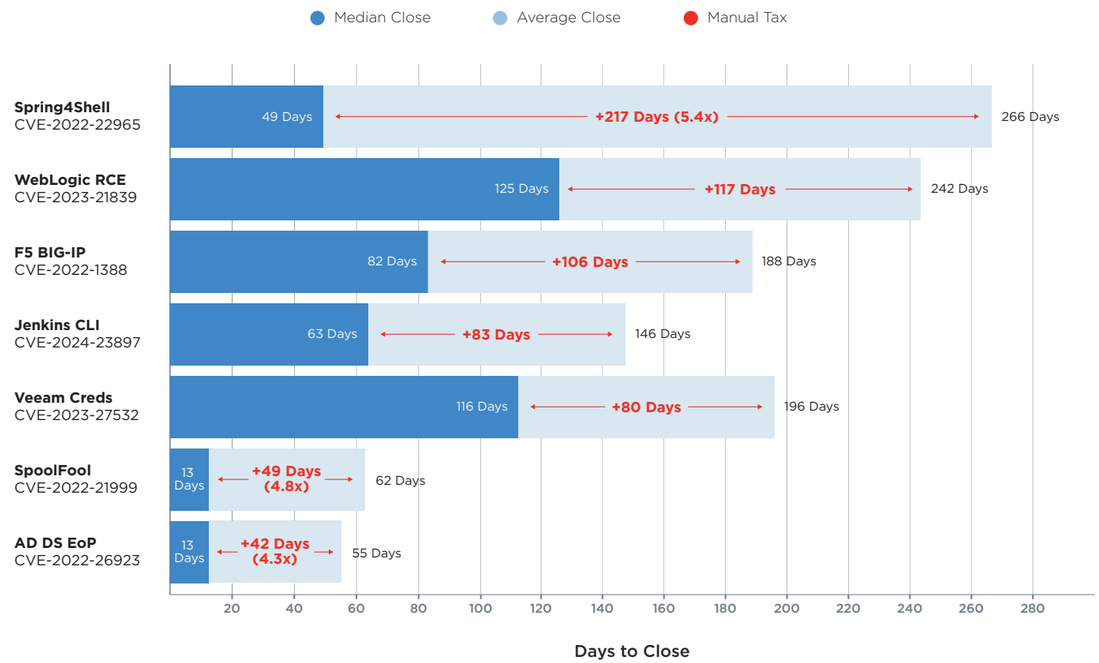

The Infrastructure Divide

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.

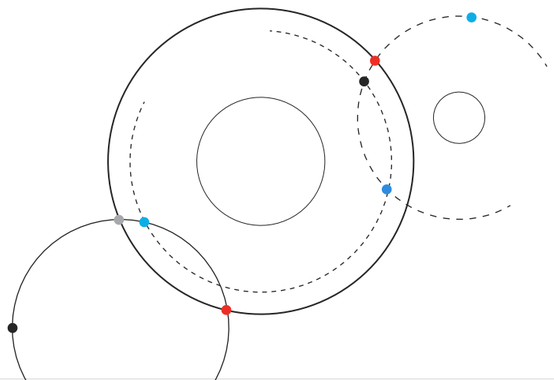

Bridging the Gap: Mitigation as a Remediation Layer

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.

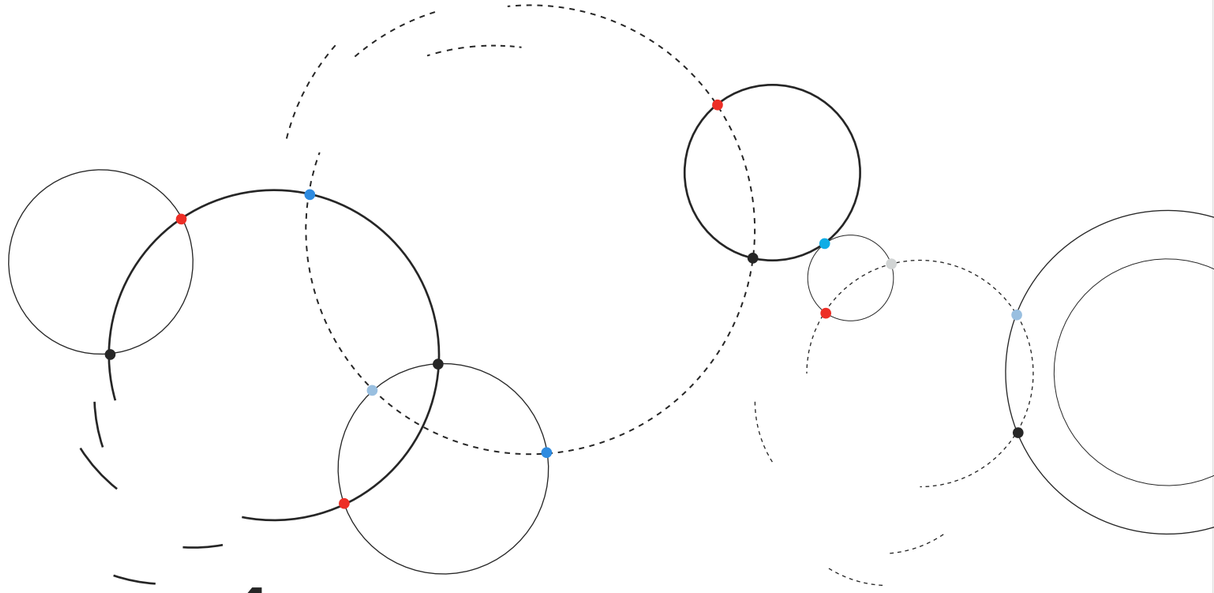

4

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.

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 Cl0p, 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.

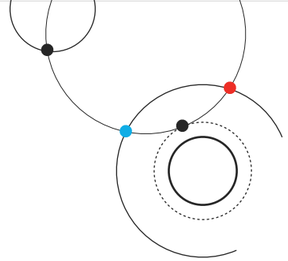

We adopt the term Risk Mass — analogous to cumulative exposure metrics long established in epidemiology and reliability engineering — to bring a dimension into vulnerability management that CVE counts and MTTR both ignore: time. The formula is straightforward — the number of vulnerable assets multiplied 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.

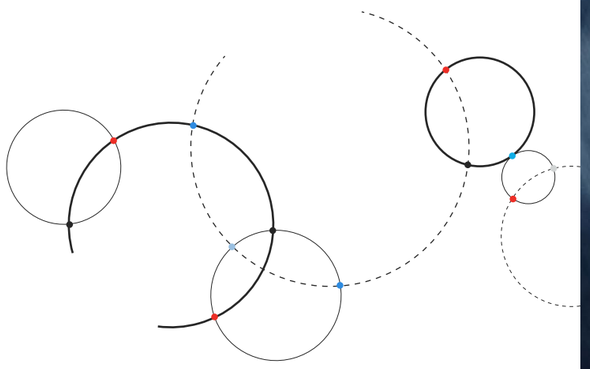

Qualys Threat Research Unit

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 (Cl0p, LockBit 3.0, Bl00dy, 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

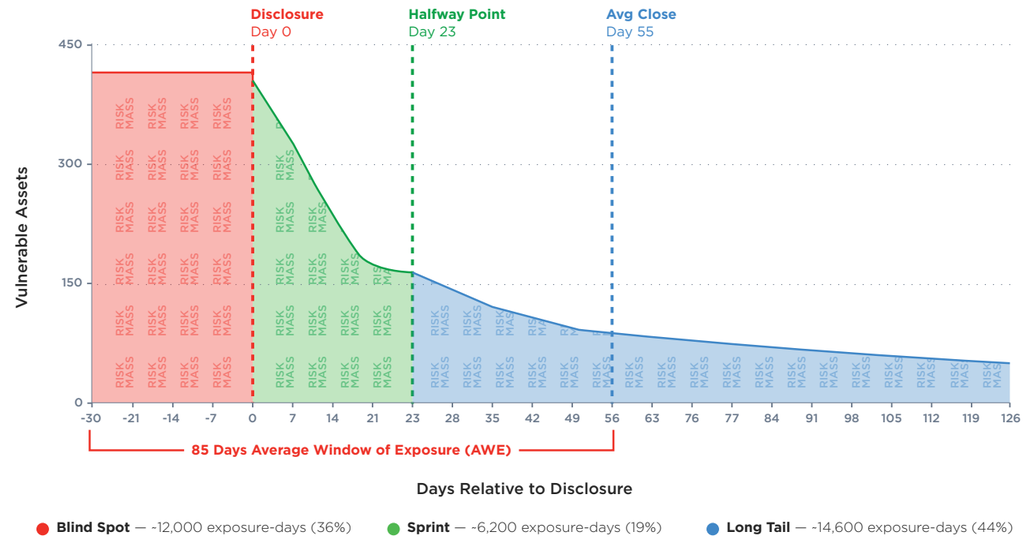

Blind Spot — ~12,000 exposure-days (36%) **Sprint** — ~6,200 exposure-days (19%) **Long Tail** — ~14,600 exposure-days (44%)

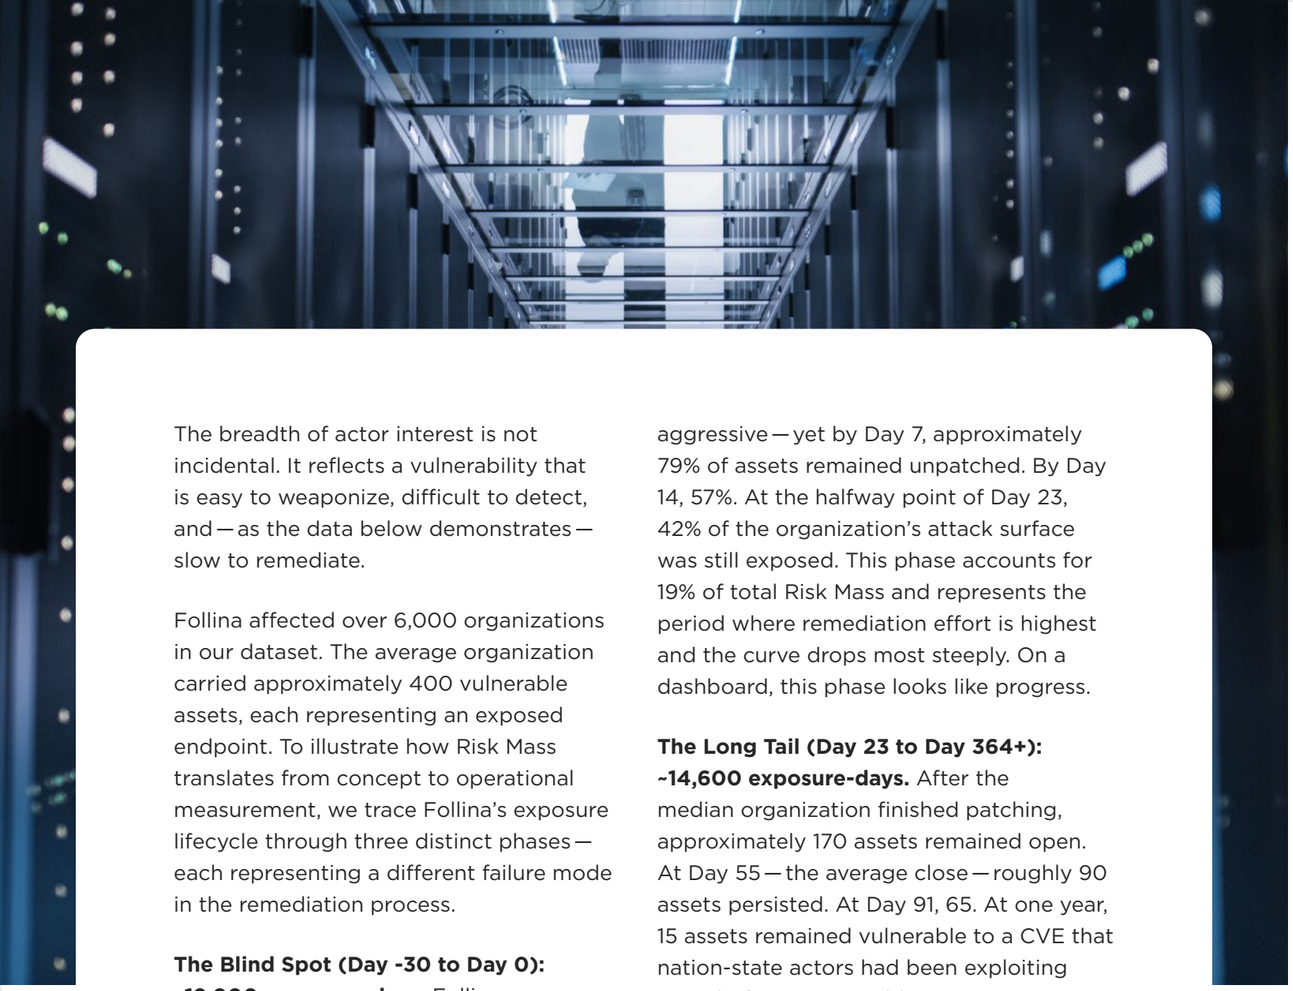

The breadth of actor interest is not incidental. It reflects a vulnerability that is easy to weaponize, difficult to detect, and—as the data below demonstrates—slow to remediate.

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.

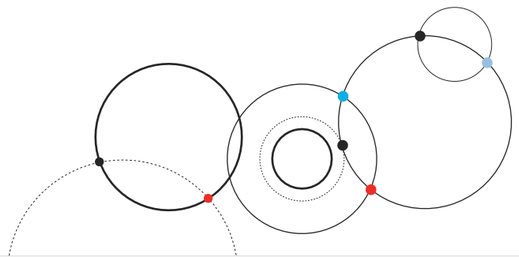

The Sub-1% Reality: In 2025, the industry disclosed **48,172** new vulnerabilities. Of those, our threat intelligence confirms that **357** — just **0.74%** — were remotely exploitable, actively weaponized in the wild, and supported by working proof-of-concept code. This **sub-1%** reality has remained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

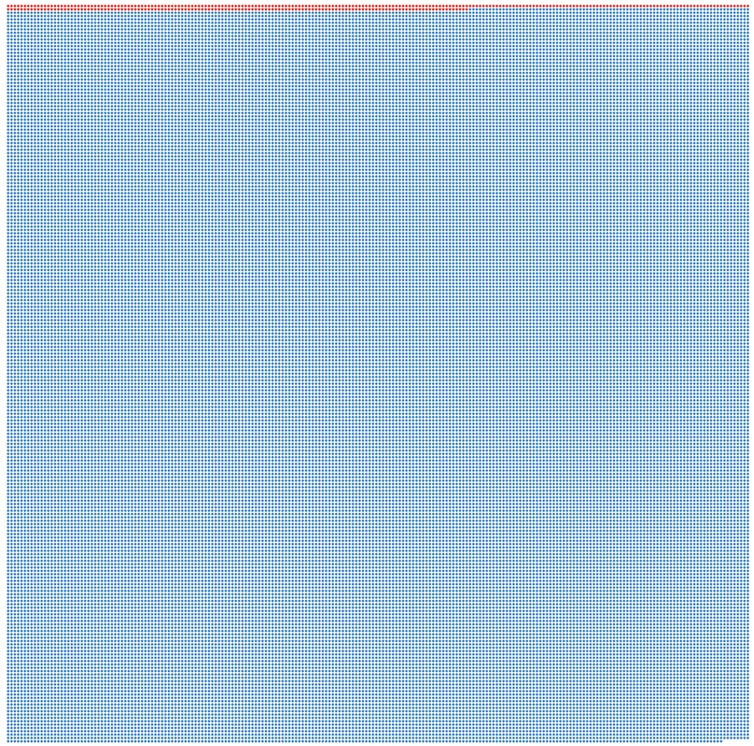

YEAR OVER YEAR

2022
25,043 disclosed
131 weaponized 0.52%

2023
28,816 disclosed
185 weaponized 0.64%

2024
39,964 disclosed
218 weaponized 0.55%

2025
48,172 disclosed
357 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.

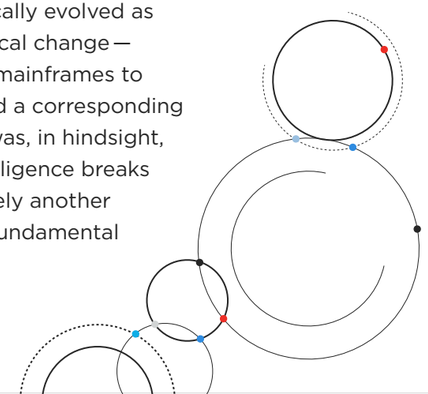

This asymmetry defines the central challenge of the current era: until defenders can match the speed and autonomy of AI-driven offense, the gap will widen. It is within this context that agentic AI orchestration emerges not as an incremental improvement but as an operational necessity. Agent Val, the 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.

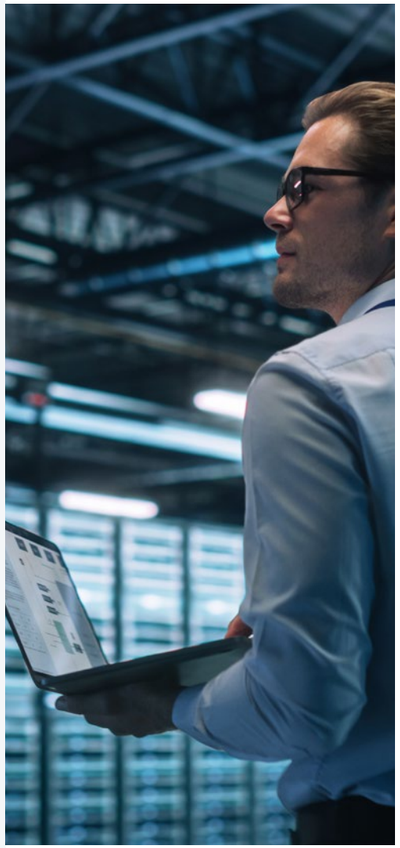

6 Operationalize or Fail

The evidence in this report — 6.5x volume growth, a TTE of -1 day, 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.

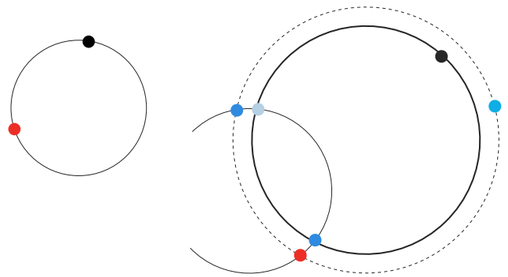

The End of Risk Whack-a-Mole

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.

Qualys®

Qualys Threat Research Unit

www.qualys.com/tru

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.

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.

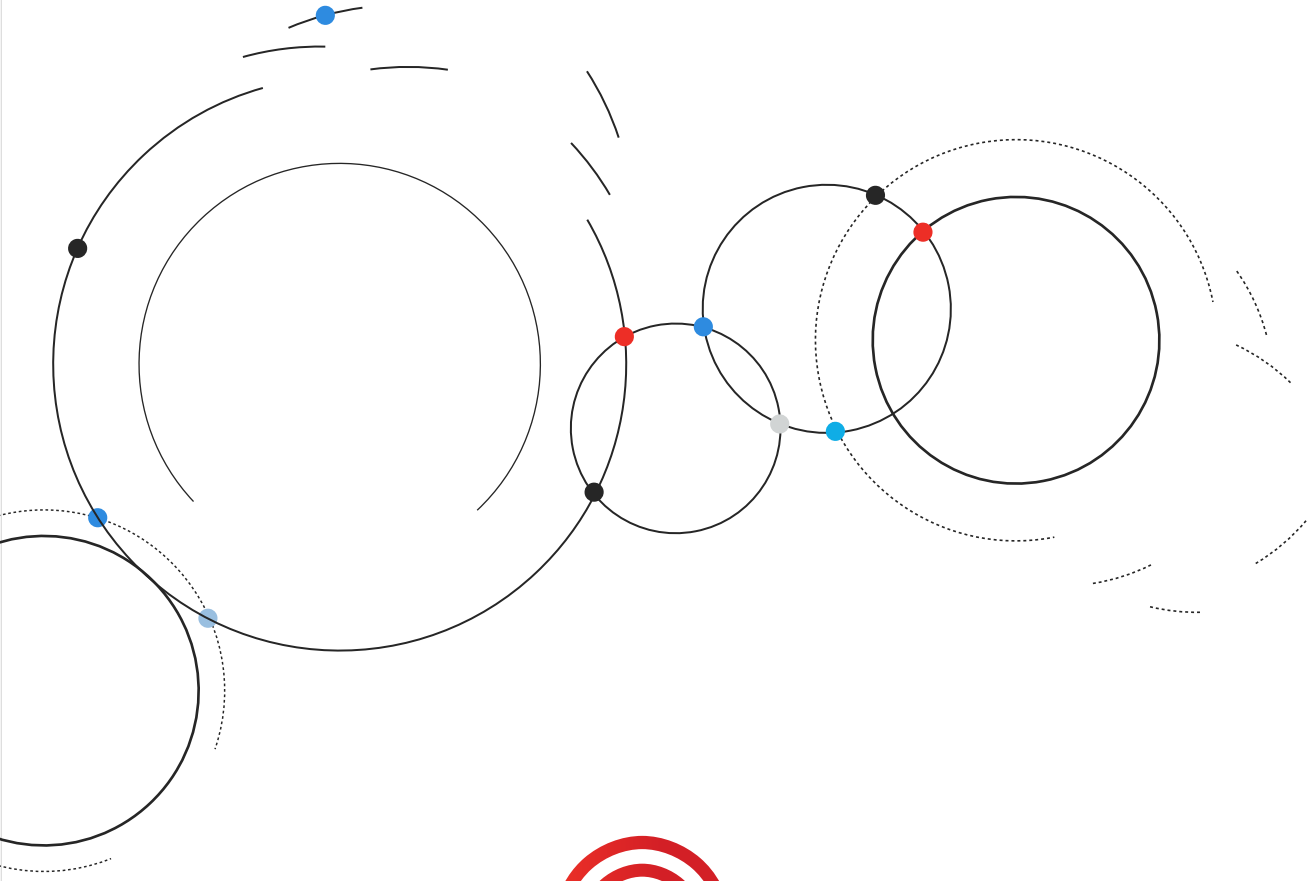

Qualys Threat Research Unit

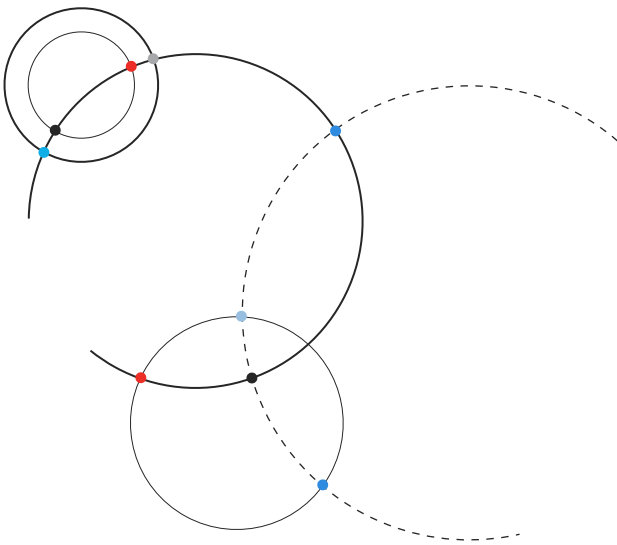

Qualys.