AGENTBOUNTY

THE REMEDIATION GAP

Why Finding Vulnerabilities Was Never the Hard Part

29

MINUTES

Attackers move

VS

32

DAYS

To patch

THE 74-DAY WINDOW

This story plays out in some version at organizations every quarter.

A financial services firm completed its annual penetration test. The results were thorough. The report was professional. It identified 47 vulnerabilities, including three rated critical — an unpatched API endpoint exposing customer records, a misconfigured IAM role granting lateral movement across cloud environments, and a legacy authentication bypass on an internal application.

The report was delivered on a Tuesday. By Thursday it had been reviewed by the security team. The following week, tickets were created in Jira. The API fix was assigned to a backend developer who was mid-sprint on a product launch. The IAM role was esc

ERROR fo
Invalid do

to the cloud team, who flagged it for their next maintenance window. The authentication bypass went to a contractor who had built the original system — he was no longer under contract.

Eleven weeks later, two of the three critical findings were still open.

Twelve weeks later, the firm disclosed a breach. The attack vector was the unpatched API endpoint — the first vulnerability on the list. The attacker had discovered it using automated reconnaissance tools three weeks after the pen test report was delivered. They had been inside the environment for 56 days before detection.

The pen test worked. Detection was not the failure point. The 74-day gap between "we know about this" and "this is actually fixed" was the failure point.

This pattern is not hypothetical. It is the norm.

THE MATH DOESN'T WORK

The security industry has become extraordinarily effective at finding vulnerabilities. Discovery tools are faster, more comprehensive, and more accessible than at any point in history. What has not kept pace is the other side of the equation — the capacity to actually fix what we find.

The data tells a consistent story across every major industry report:

The average time to remediate an internet-facing host or cloud vulnerability is 61 days.[1] The average time from public disclosure to active exploitation is 5 days.[2] That is not a gap — it is a forfeit. For every infrastructure vulnerability sitting in a remediation queue, attackers have roughly a 12:1 time advantage.

In some cases, even that ratio is generous. The 2025 Verizon Data Breach Investigations Report — the industry's largest breach analysis, covering 22,052 incidents and 12,195 confirmed breaches across 139 countries — found that for edge device vulnerabilities, the median time to mass exploitation was zero days.[3] The vulnerabilities were being

actively exploited before organizations even knew they needed to patch. The median time to remediate those same vulnerabilities was 32 days. And only 54% were fully remediated at all.[3]

This is not a theoretical risk. Exploitation of edge devices and VPNs increased eightfold in a single year, jumping from 3% to 22% of all breaches.[3] Vulnerability exploitation as an initial access vector grew 34% year over year and now accounts for 20% of all confirmed breaches — approaching the frequency of stolen credentials.[3] Ransomware was present in 44% of breaches, up from 32%.[3] Attackers are not waiting for defenders to catch up. They are accelerating.

Notably, MTTR for critical web application vulnerabilities has improved — dropping to 35 days, down significantly from prior years.[1] Application security tooling is getting better. But infrastructure and cloud findings still average nearly twice as long to remediate. The gap is not closing uniformly. It is migrating — away from application code and deeper into the infrastructure, identity, and cloud layers where automated code fixes cannot reach.

Large enterprises leave 45.4% of discovered vulnerabilities unresolved after 12 months, with 17.4% of those classified as high or critical severity.[1] These are not obscure, low-severity findings buried in scan noise. Nearly half of all identified vulnerabilities — including those discovered through professional penetration testing — remain open a full year later. In the network and device layer, the numbers are even worse.

The median time-to-exploit after public disclosure has collapsed to 5 days.[2] Twelve percent of disclosed vulnerabilities are weaponized within 24 hours and 29% within a week. In 2025, 884 CVEs were identified with first-time exploitation evidence — and 28.96% of those were exploited on or before the day the CVE was even published.[4] Over 48,000 new CVEs were disclosed in 2025 alone, up from roughly 40,000 the prior year.[4] The window between "this vulnerability exists" and "someone built an exploit for it" is measured in hours. The window between "we found it" and "we fixed it" is still measured in months.

241

DAYS

Mean time to identify & contain a breach — lowest in 9 years

45.4%

UNRESOLVED

Vulnerabilities still open after 12 months at large enterprises

48K+

NEW CVES

Disclosed in 2025 alone — up from ~40K the prior year

Increasing test frequency improves outcomes — but reveals the real bottleneck. One midsize healthcare firm that transitioned from annual to quarterly pen testing reduced unresolved vulnerabilities by 42% within six months.[5] That's meaningful, but it also reveals the deeper problem: more frequent reporting creates more frequent pressure to remediate. The testing itself was never the bottleneck.

The cumulative picture is damning. The industry spends billions on detection, scanning, testing, and monitoring. The result is an ever-growing backlog of known vulnerabilities that organizations know about, have documented, and still cannot fix fast enough. The bottleneck was never discovery. The bottleneck is remediation.

ATTACKERS ARE GETTING FASTER. DEFENDERS ARE NOT.

The remediation gap would be concerning even if the threat landscape were static. It is not. Attacker speed is accelerating at a rate that fundamentally changes the math for defenders.

The CrowdStrike 2026 Global Threat Report — released in February 2026 and covering adversary activity throughout 2025 — found that the average eCrime breakout time

dropped to 29 minutes.[6] That is the interval between an attacker's initial compromise and their first lateral movement to another system in the target network. It represents a 65% increase in speed from the prior year. The fastest observed breakout was 27 seconds.[6] In one intrusion, data exfiltration began within four minutes of initial access.[6]

Consider what this means against the remediation data. Organizations take 32 to 61 days to remediate a known vulnerability. Attackers need 29 minutes — or less — to move from initial access to full network compromise. The gap between defender remediation speed and attacker operational speed is not narrowing. It is widening.

The acceleration is being driven by several converging factors. Eighty-two percent of CrowdStrike's 2025 detections were malware-free — attackers are using valid credentials, trusted identity flows, and approved SaaS integrations to move through environments without triggering traditional endpoint security tools.[6] Cloud-conscious intrusions rose 37%, with state-nexus cloud attacks surging 266%.[6] Zero-day exploitation increased 42% year over year, with adversaries weaponizing vulnerabilities before public disclosure.[6] AI-enabled adversaries increased their attack volume by 89%.[6]

ATTACKER ACCELERATION — 2025 OBSERVED BEHAVIORS

82%
Malware-free detections

+37%
Cloud intrusions

+42%
Zero-day exploitation

+89%
AI-enabled attack volume

Source: CrowdStrike 2026 Global Threat Report[6]

These are not theoretical projections. These are observed behaviors from the past 12 months. The organizations sitting at Level 0 or Level 1 remediation maturity — and that is most organizations — are operating on a timeline that assumes attackers move in days or weeks. Attackers are now moving in minutes.

THE WORKFORCE CANNOT CLOSE THE GAP

The natural response to a remediation bottleneck is to hire more people. The data makes clear why that strategy has failed.

The ISC2 2025 Cybersecurity Workforce Study — the largest annual survey of the cybersecurity profession, covering 16,029 respondents globally — found that 88% of organizations have experienced at least one significant cybersecurity consequence directly attributable to a skills shortage.[7] Sixty-nine percent have experienced more than one.[7]

The problem is no longer just headcount. For the first time, ISC2 declined to publish a global workforce gap estimate, noting that respondents now prioritize the need for critical skills over the need for more bodies.[7] Fifty-nine percent of respondents cited critical or significant skill shortages within their teams — a 15-point increase from the prior year.[7] Thirty-three percent said their organizations do not have the resources to adequately staff their teams. Twenty-nine percent said they cannot afford to hire people with the skills they actually need.[7]

The consequences show up in operations. Twenty-six percent of organizations reported oversights in cybersecurity processes due to skills gaps. Twenty-five percent have been forced to place underqualified personnel into roles to cover shortfalls. Twenty-four percent have misconfigured systems as a direct result of insufficient expertise.[7]

And the people who remain are burning out. Forty-eight percent of respondents said they feel exhausted from trying to stay current on the latest threats and emerging technologies. Forty-seven percent feel overwhelmed by workload.[7] Seventy-two percent believe that reducing security staff significantly increases the risk of a breach — yet budget cuts (36%) and layoffs (24%) continue, even as they show signs of leveling off.[7]

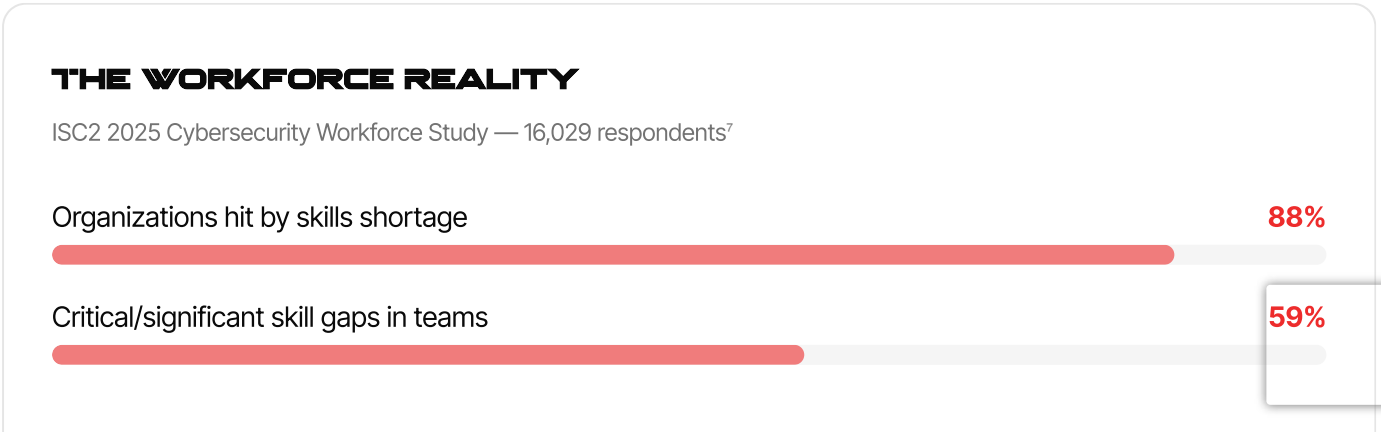

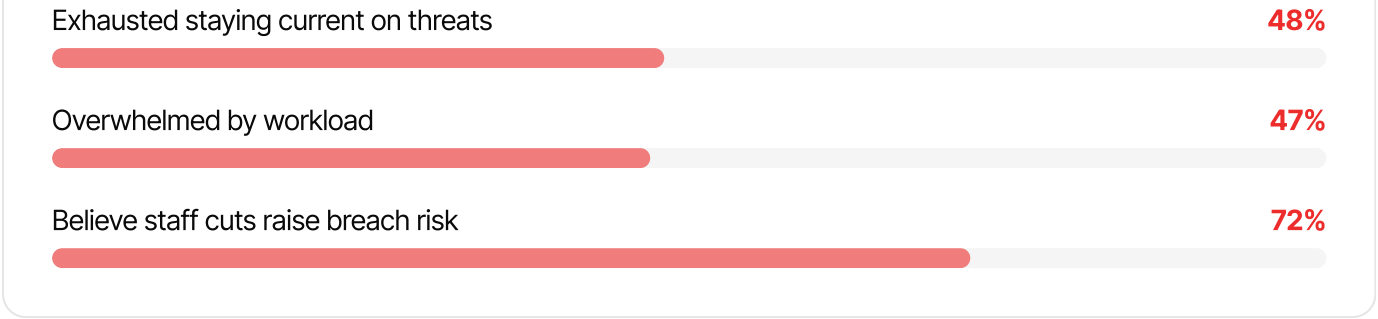

You cannot hire your way out of a remediation gap when the skills do not exist in sufficient quantity, the people you have are overwhelmed, and the budget to acquire more is constrained. The only path forward is to change the ratio between human effort and remediation output. That means automation — not as a buzzword, but as an architectural requirement.

THE COST OF THE GAP

The remediation gap is not an abstract operational inefficiency. It has a price.

The IBM Cost of a Data Breach Report 2025 — based on 600 organizations across 17 industries and 16 countries, covering breaches from March 2024 through February 2025 — puts the global average cost of a data breach at $4.4 million.[8] In the United States, that figure rises to $10.22 million. In healthcare, it reaches $7.42 million.[8]

But the more consequential finding is the cost differential between organizations that have deployed AI and automation extensively in their security operations and those that have not. Organizations with extensive AI and automation reported average breach costs of $3.62 million. Those without: $5.52 million.[8] That is a $1.9 million gap — and the automated organizations also reduced their breach lifecycle by an average of 80 days.[8]

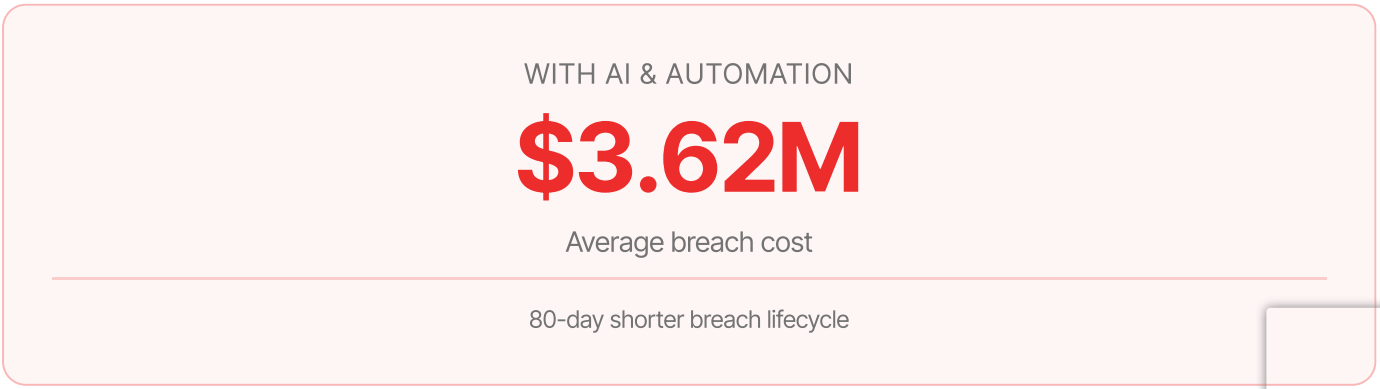

WITHOUT AI & AUTOMATION

$5.52M

Average breach cost

$1.9M more per breach

The mean time to identify and contain a breach dropped to 241 days in 2025 — the lowest in nine years — driven primarily by organizations that invested in automation.[8] Among organizations that experienced a breach and recovered, most took more than 100 days to fully recover.[8] Nearly half planned to raise the prices of their goods and services as a result of the breach.[8]

The pattern is consistent across every data point: speed of remediation directly correlates to cost of breach. Organizations that fix faster pay less. Organizations that automate remediation fix faster. The financial case for closing the remediation gap is not speculative. It is documented, quantified, and published annually by the most recognized research in the industry.

A FRAMEWORK FOR EVALUATING WHERE YOU ACTUALLY STAND

Not all remediation is created equal. The market is filled with claims of "automated remediation," but the term has been stretched to cover everything from generating a PDF report to autonomously deploying a verified fix. To evaluate where any organization or vendor actually sits, it helps to define a clear maturity spectrum.

REMEDIATION MATURITY FRAMEWORK

Where does your organization sit?

Level 0: Report and Pray

Manual PDF reports, no tracking

Level 1: Prioritized Guidance

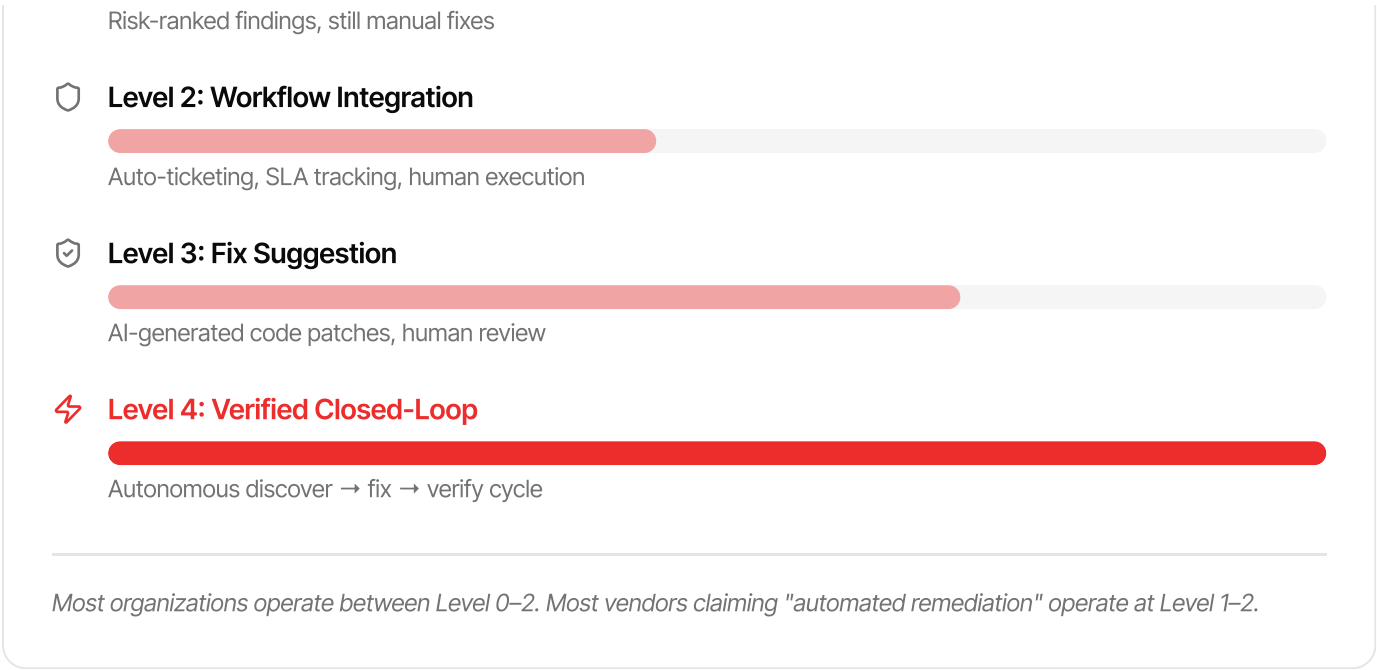

LEVEL 0: REPORT AND PRAY

A penetration test is conducted. A report is delivered — typically a PDF, often exceeding 100 pages. The findings are organized by severity. Remediation guidance is included in general terms: "patch the affected system," "restrict access to the endpoint," "update the dependency."

The report is emailed to the security team. The security team reviews it, creates tickets, and assigns them to engineering. Engineering prioritizes the tickets against their existing sprint backlog. Weeks pass. Some findings are addressed. Many are not. The next annual test reveals that 40–60% of previous findings are still open.

This is where the majority of the industry still operates. The test works. The report is accurate. The remediation process is entirely manual, entirely dependent on human bandwidth, and structurally unable to keep pace with the rate of new findings.

LEVEL 1: PRIORITIZED GUIDANCE

The platform goes beyond a static report and provides risk-based prioritization. Findings are ranked not just by CVSS score but by exploitability, asset criticality, and business context. The security team sees which vulnerabilities matter most and can allocate resources accordingly.

This is a genuine improvement over Level 0. Knowing what to fix first reduces wasted effort and focuses attention on the highest-risk items. But the remediation itself is still entirely manual. The platform tells you what to fix and why it matters. You still have to do all the work.

Most breach-and-attack simulation platforms, continuous validation tools, and advanced pen testing platforms operate at this level. They have improved the intelligence of the finding. They have not touched the mechanics of the fix.

LEVEL 2: WORKFLOW INTEGRATION

Findings automatically generate tickets in Jira, ServiceNow, or equivalent systems. Remediation playbooks are attached. Assignments route to the appropriate team based on asset ownership. SLAs are tracked. Escalation paths are defined.

This level reduces the administrative overhead of remediation and creates accountability through workflow automation. The fix process is faster because the handoff is smoother. But the actual remediation — writing the patch, changing the configuration, deploying the update — still depends on a human engineer reading the ticket, understanding the context, developing the fix, testing it, and deploying it.

The bottleneck has shifted from "we didn't know what to fix" to "we don't have enough engineers to fix everything fast enough." The workflow is optimized. The fundamental constraint — human execution speed — remains.

LEVEL 3: FIX SUGGESTION

The platform generates a recommended fix. For application-layer vulnerabilities, this may take the form of a code suggestion or a draft pull request. For infrastructure findings, it may produce a configuration recommendation or a policy update. A human reviews the suggestion, validates it, and implements it.

This is where a small number of platforms have begun to operate, primarily in the application security and DevSecOps space. AI-powered code analysis can examine a

vulnerable function, understand the context, and propose a patch. The human developer reviews the suggestion and merges it if appropriate.

This represents meaningful progress. The time from finding to fix decreases significantly because the engineer is reviewing a proposed solution rather than building one from scratch. However, the scope is limited. These suggestions work for known vulnerability patterns in application code — dependency updates, injection fixes, configuration corrections. They do not extend to network-level findings, identity and access management gaps, infrastructure misconfigurations that require architectural changes, or multi-step attack paths that span multiple systems and layers.

The fix is suggested, not executed. The scope is application-layer, not enterprise-wide. And verification — confirming the fix actually resolved the vulnerability — still requires a separate testing cycle.

LEVEL 4: VERIFIED CLOSED-LOOP REMEDIATION

The system discovers the vulnerability through active offensive testing. It validates exploitability by actually exploiting it in a controlled manner. It generates the specific fix — not a suggestion or a recommendation, but the actual patch, configuration change, or policy update required. It presents the fix for human approval. Upon approval, the fix is applied. The system then retests the specific vulnerability to confirm the fix is effective. The loop is closed.

The human role shifts from executor to approver. The system handles discovery, exploitation, fix generation, and verification. The human provides oversight, judgment, and authorization. This preserves the trust and governance that enterprise security requires while removing the execution bottleneck that causes remediation delays.

At this level, the scope extends beyond application code to encompass infrastructure, network, identity, and cloud configuration findings. A vulnerability in an IAM policy is not just flagged — the corrected policy is generated. A misconfigured firewall rule is not just reported — the corrected rule is drafted. A multi-step attack path that chains a weak password policy with excessive role permissions and an unpatched service is not just documented — each component of the chain receives a targeted fix.

This is where remediation stops being a process and becomes a capability.

The shift from reactive remediation to continuous, closed-loop security is not a fringe idea. Gartner predicts that by 2026, organizations that prioritize their security investments based on a Continuous Threat Exposure Management (CTEM) program will be three times less likely to suffer a breach.[9] The analyst consensus is clear: continuous discovery, validation, and remediation — operating as an unbroken cycle — is the direction the industry must move. The question is not whether this shift will happen. It is whether your organization will make it before attackers exploit the delay.

Where does your organization sit? Most security teams will recognize themselves somewhere between Level 0 and Level 2. Most vendors claiming "automated remediation" operate at Level 1 or Level 2, with a small number reaching Level 3 for application-specific findings. Level 4 requires a fundamentally different architecture — one that integrates offensive testing, exploit validation, fix generation, and verification into a single continuous system.

CODE FIXES ARE NECESSARY. THEY ARE NOT SUFFICIENT.

A new category of platform has emerged that deserves recognition. Several vendors can now generate code-level fixes for application vulnerabilities detected through scanning — automatically creating pull requests that address known CVEs, dependency issues, injection vulnerabilities, and infrastructure-as-code misconfigurations. For development teams drowning in security backlogs, this is a genuine step forward. It reduces MTTR for application findings and keeps security from blocking release velocity.

But enterprise security does not begin and end in the code repository.

When a penetration test reveals that an attacker can chain a phishing vector with a misconfigured VPN, escalate privileges through a poorly scoped Active Directory group policy, and exfiltrate data through an overly permissive cloud storage bucket — no pull request fixes that. The vulnerability exists across identity systems, network architecture, cloud infrastructure, and human processes. The remediation requires configuration

changes across multiple platforms, policy updates in identity providers, network segmentation adjustments, and potentially architectural redesign.

Application-layer auto-fix addresses one dimension of a multidimensional problem. It is Level 3 remediation scoped to a single attack surface. For the enterprise running a hybrid environment with cloud workloads, on-premises infrastructure, third-party integrations, and a distributed workforce, the remediation gap exists precisely in the spaces between application code — in the infrastructure, identity, network, and operational layers where code-level fixes do not reach.

Closing the remediation gap requires a system that operates across the full enterprise attack surface, not just the software development lifecycle. It requires remediation that is as comprehensive as the attack itself.

ENDNOTES

1. Edgescan, 2025 Vulnerability Statistics Report, 10th Edition.

2. Google Mandiant, M-Trends 2025, 16th Edition.

3. Verizon, 2025 Data Breach Investigations Report. 22,052 incidents, 12,195 confirmed breaches, 139 countries.

4. VulnCheck, State of Exploitation 2026. 884 KEVs with first-time exploitation evidence during 2025.

5. Case study cited in Deepstrike.io (2025). Single organization; not an industry-wide benchmark.

6. CrowdStrike, 2026 Global Threat Report (February 2026). 281+ tracked threat groups.

7. ISC2, 2025 Cybersecurity Workforce Study (December 2025). 16,029 respondents globally.

8. IBM Security / Ponemon Institute, Cost of a Data Breach Report 2025 (July 2025). 600 organizations, 17 industries, 16 countries.

9. Gartner, "Implement a Continuous Threat Exposure Management (CTEM) Program," October 2023.

ABOUT AGENT BOUNTY

This paper was produced by the Agent Bounty research team. Agent Bounty is a continuous security platform designed to operate at Level 4 of the remediation maturity framework outlined in this paper. The platform integrates offensive testing, exploit validation, fix generation, and post-remediation verification into a single closed-loop system — spanning application, infrastructure, identity, and cloud attack surfaces.