astra

ANNUAL REPORT

State of Continuous Pentesting Report 2026

Built in the trenches.
Backed by data.
For every security decision.

PUBLISHED BY
Astra Security

www.getastra.com

Foreword

Security in 2026 is an architecture problem.

When we started the year, we kept seeing the same gap: engineering teams shipping daily, cloud infrastructure changing weekly, and AI features landing in production without a security review. The attack surface was moving continuously. The security programme was measuring it annually.

This report is our attempt to close that gap with data, backed by 6.8 million findings from 150K+ scans and 8K+ pentest engagements across web, API, cloud, mobile, and network infrastructure.

> *"The uncomfortable truth, as we looked at this data, is that most organisations are under-secured because they are mis-measured."*

The budget, intent, and tools all exist, but the metric being measured (total vulnerability count, monthly scan volume, CVE patch rate) has structurally decoupled from actual risk. In 2025, critical vulnerabilities grew at 14.6x the rate of everything else, and most dashboards registered it as just more volume.

That gap between measurement and reality is where breaches live. Not because the numbers are alarming, though some of them are, but because the gap is now wide enough to make confident, well-funded security programmes actively misleading. A CISO walking into a board review with a clean total-count dashboard in Q4 2025 was presenting last season's weather as tomorrow's forecast.

We built Astra because we believed continuous, evidence-based security was the only honest answer to a continuously changing attack surface. What the 2025 data tells us is that the case for that is no longer philosophical.

- Cloud overtook the web as the primary attack surface in three separate quarters.
- The highest-value cloud exposures were found in mobile apps by analysts who weren't looking for them.
- A new vulnerability class arrived in production with no CVE, patches, or remediation playbook.

Simply put, the data in this report won't make for comfortable reading. Maybe that's a good thing, since we are up for the most uncomfortable year yet.

Shikhil and Ananda

FOUNDERS, ASTRA SECURITY

Introduction

> 2026 is already a quarter old. How's it looking so far?

If your security programme is running on last year's assumptions, **the answer is: more complicated than your dashboard suggests.**

Here's what actually caught us off guard: the most dangerous month of 2025 wasn't the one with the most vulnerabilities. The biggest cloud exposures weren't found by cloud engineers. A vulnerability class that didn't exist in your triage queue in 2024 is already in production with no CVE, no vendor patch, and no established playbook for what to do when you find it.

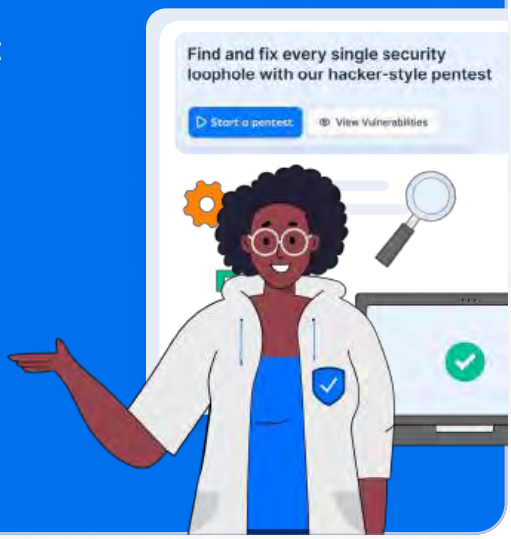

That gap, between what happened and what most programmes were set up to see, is what this report (with **6.8 million findings, 150,000+ scans, and 8,000+ engagements**) is about.

At the end of it, we will ask you three questions: not about your budget or tooling, but about what your programme can actually see. Each has a clear, data-backed answer that any properly instrumented programme should be able to produce. Most organisations can answer one confidently, fewer attempt two, and almost none can answer all three.

Most are surprised by which one.

The rest of 2026 is still yours to shape; this is what you're shaping it against.

Executive Summary

THE METRIC YOU'RE USING WRONG

275%

Total vulnerability growth in 2025.

14.6x
How fast critical vulnerabilities grew relative to everything else.

1 in 10
Findings in 2025 were critical. In 2024, it was 1 in 40.

THE AUTONOMOUS PENTESTING SHIFT

80x
Faster testing with autonomous pentesting (minutes, not weeks)

173
OWASP APTS requirements for autonomous governance

4 Autonomy levels
Defined in OWASP APTS (L1 through L4)

THE 30 DAY BLIND SPOT

30 days

How far ahead does your scan data actually predict (not the same month)?

1.8M
Vulnerabilities were found in December alone, preceded by the quietest scanning month of the year.

CLOUD: THE TROJAN HORSE

44x GROWTH

Cloud vulnerability growth in a single year.

39%
Of all findings came from the cloud, while manual pentest coverage grew only 1.23x

3
Separate quarters where cloud overtook web as the primary attack surface.

THE 37-TO-1 PROBLEM

37×
Cloud vulnerability growth outpaced cloud testing growth by this factor.

7,480
Average vulnerabilities found per cloud pentest engagement.

THE LARGEST ATTACK SURFACE HAS NO PRICE TAG

80%

Of S3 and AWS credential exposure was found inside mobile apps.

2.4x
The yield of a cloud pentest compared to a web pentest.

THE IDOR PLAGUE

6 of 6
Tested surfaces where IDOR appears simultaneously.

$1.1M
Tracked financial exposure from IDOR alone, the highest of any vulnerability class.

YOU CAN'T PATCH YOUR WAY OUT OF A DESIGN FLAW

91%
Drop in CVE- tracked disclosures on the Astra Security platform.

20x
Manual vulnerability discovery growth in the same period.

48 seconds
How often was a critical vulnerability found in 2025?

THE 2025 DATASET

$334

Average tracked financial exposure per vulnerability

$2.37B

Total tracked financial exposure

ORGANIZATIONS THE DATA CANNOT SEE

22%
Of organisations ran one engagement and went dark.

29%
Of organisations test a single surface only.

5%
Manufacturing and energy share of testing customers: the most targeted sectors by nation-state actors

AI INTRODUCED A VULNERABILITY CLASS THAT DIDN'T EXIST IN 2024

2
AI vulnerability classes are appearing in production pentests for the first time.

$35K
Total tracked financial exposure from AI vulnerability classes in 2025, the floor, not the ceiling.

Q4 HAD THE NUMBERS, Q3 HAD THE RISK

29%
How much more dangerous was Q3, per the findings than Q4?

63%
More raw findings in Q4 (the loud quarter, not the dangerous one).

The Metric You're *Using Wrong*

> The Problem With Most Security Reports Is Not The Data But The Question They Answer.

Most reports (penetration testing and cybersecurity trend analysis alike) answer the question: how many vulnerabilities were found last year? That question is easy to measure and, to be blunt, increasingly useless.

The 2026 Planning Assumption

In 2025, Astra Security's Platform discovered 6.8 million vulnerabilities, a 275% increase from the previous year (reflecting both platform adoption growth and broader expansion of the attack surface under coverage).

While the 275% figure (growth in both threat landscape and testing coverage) in itself is enough to give most CTOs a good nightmare or two, it presents a fatally incomplete picture, because severity in 2025 is moving in opposite directions.

275%

TOTAL VULNERABILITY GROWTH

Forecast models point to 2.7x (minimum) growth in 2026, treat it as a direction, not a guarantee.

14.6x

CRITICAL SEVERITY GROWTH

If the trend holds, 1 in 5 findings by the end of 2026 will be critical.

+7PP

CRITICAL SHARE 2024 TO 2025

The floor has shifted and won't shift back without genuine architectural changes.

The total volume grew by 275%, but the critical vulnerabilities alone grew nearly 4x, as low-criticality vulnerabilities rose by 1.5x. Both are inside that headline number, weighted together, indistinguishable to any dashboard that counts without sorting.

An organisation that sizes its 2026 response to the total count will overfund low-severity

remediation and underfund the critical queue (i.e., patching CSS while the authentication layer is on fire) repeating the story we saw in 2025.

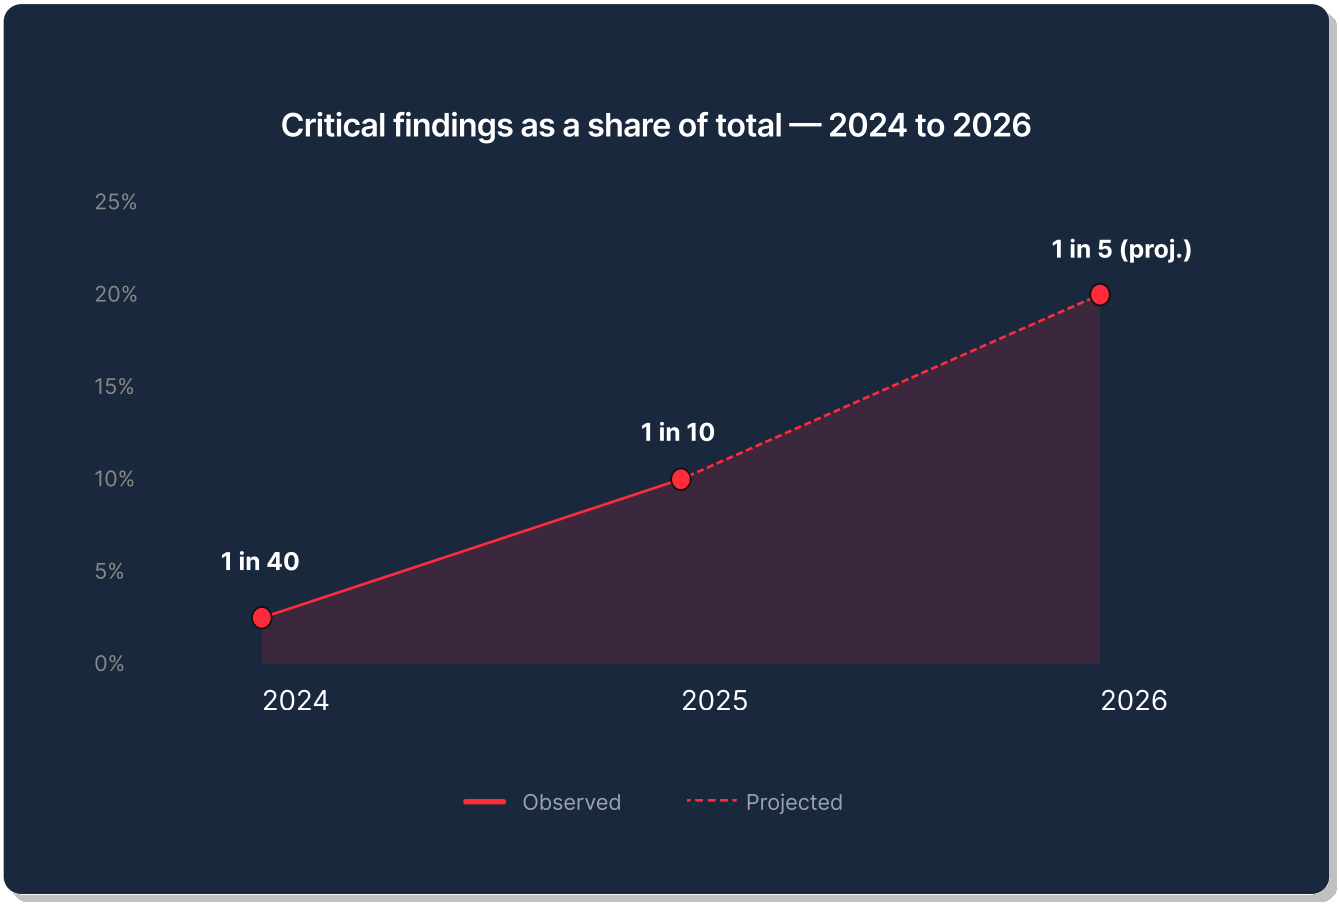

Simply put, 1 in every 10 findings last year was a critical vulnerability. In 2024, this number was closer to 1 in 40. If the composition trend continues, **by the end of 2026, 1 in 5 vulnerabilities will be critical.**

Your traditional dashboards will register this as just more volume, whereas a severity-weighted dashboard would map it to a fundamentally different threat environment.

What is truly worrying, most organisations are running the former and believing it shows them the latter.

The 2025 trend line points to 2.7x growth in 2026, though every structural driver behind that number, cloud expansion, critical severity growth, and new attack surfaces entering scope for the first time, suggests the real number is more likely to exceed that estimate than fall short of it.

Worse, the practical consequences of this mean that a 2026 security investment calibrated against a 2025 risk profile and total count would **underspend on the real threats** while overspending on low-severity, low-impact remediation cycles that moved the count but kept your risk profile score stagnant at a strong D.

Your risk trajectory in 2026 will primarily depend on severity-weighted critical exposure and NOT total count, as it has till now

Severity Breakdown: The Floor Is Rising

SEVERITY	2025 SHARE	2024 SHARE	SHARE SHIFT	COUNT GROWTH
Critical	**9.50%**	2.30%	**+7.2pp**	1360%
High	**8.50%**	7.20%	**+1.3pp**	315%
Medium	34.30%	35.10%	-0.8pp	243%
Low	28.80%	40.60%	-11.8pp	149%
Info	19.00%	14.80%	+4.2pp	352%
Total	19.00%	14.80%	+4.2pp	275%

Table 1: Every vulnerability added to the count in 2025 is, on average, more dangerous than one added in 2024. The floor is rising.

In 2026, plan for the severity floor and frequency of critical and high vulns instead of the totals.

Surge Patterns & What They Mean?

While historical data from 2025 is far from predictive, it can still be used as a template for 2026. Growth in vulnerabilities was episodic rather than linear, driven by infrastructure expansion events — new cloud environments, and new asset types coming under coverage for the first time — all of which are only accelerating.

January 2025 opened with 600K findings (6x of the same month the previous year), 700K+ in September, as December closed at 1.8M in a single month (nearly 65% of 2024 total). These seeming anomalies and outliers were actually the predictable output of cloud infrastructure expanding faster than a wildfire (or gossip) and of partially sporadic testing.

In 2026, that lag will only rise, i.e., more months that look like December 2025, longer vulnerability lists, and effective remediation cycles.

JAN 2025

600K

BASELINE

SEPT 2025

700K

+16.6% GROWTH

DEC 2025

1.8M

CRITICAL PEAK

An organization relying on monthly trend lines as a risk proxy in 2026 will cycle between false alarms and false calm, with each spike interpreted as a crisis and every quiet month as progress. Unfortunately, neither would be accurate.

This chart shows what 2025 looked like; 2026, at a higher amplitude, would show higher and more frequent peaks.

In 2026, the spikes will look identical. The only variable is whether your team treats them as surprises or scheduled events.

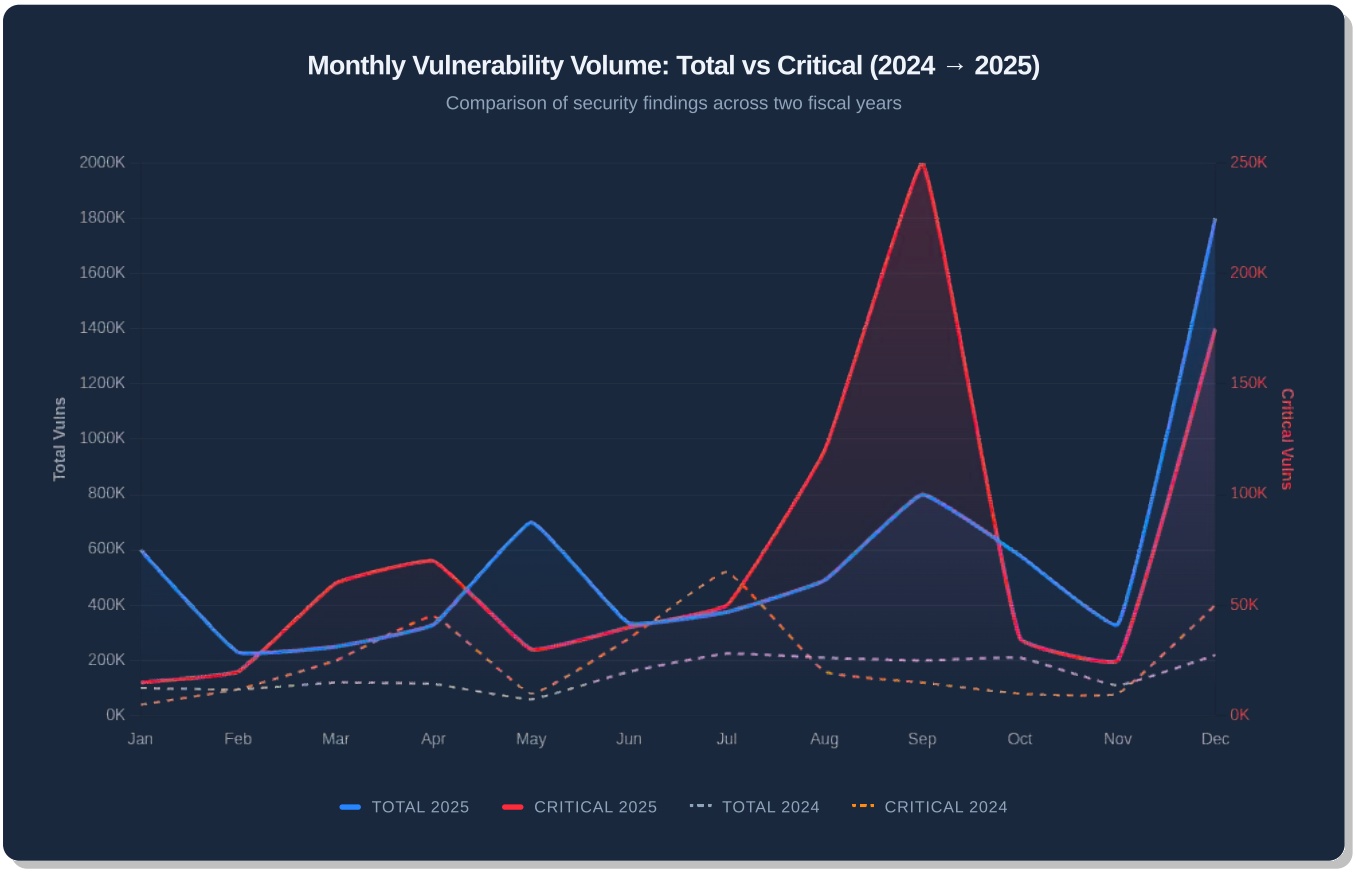

Fig.1: Monthly vulnerability volume 2024 vs 2025. December 2025 alone (1.8M) exceeds the full 2024 annual total. The 2026 baseline will likely open where December 2025 left off.

The 2026 Triage Crisis

If critical findings continue to grow at a 5x rate relative to the overall volume, the signal-to-noise ratio will compound each quarter. A team sized to the total count in 2025 is already under-resourced for the Critical queue; in 2026, it may be fatally so.

Manual pentest findings, where human judgment filters noise before a finding is logged, grew 19.7x. While many would like to chalk it up to inflation and our market share (which contributed 11.5% to this increase), it far more importantly reflects a genuine expansion of the test surface into cloud, API, and network infrastructures that previously often went untested or yielded only surface-level findings.

In 2026, this surface level is expected to be larger still. The organizations that moved to continuous pentesting coverage in 2025 are already ahead, while the others heading into 2026 with a triage team sized for 2024's threat profile are in for a critical queue that has quadrupled since the job descriptions were written.

The metric the data argues for, and the only one that will remain useful in 2026, does not count. It is severity-weighted exposure: how many critical and high findings are open, across what attack surface, at what velocity. That number grew by an order of magnitude in 2025, while the headline count grew by 3 times.

The gap between those two growth rates will widen further. Organizations still reporting lower numbers are not just measuring wrong; they are actively creating a blind spot in the infrastructure and fabric of their security programme.

What does this mean for how you build in 2026?

Before 2026 planning is finalised, replace the total vulnerability count with three numbers in your reporting stack: open Criticals, open Highs, and the severity-weighted exposure rate for your highest-risk asset type.

Everything else is a lagging indicator that will tell you about the crisis after it's already in your backlog

The 30 Day Blind Spot

> December 2025 saw more vulnerabilities than the entire year of 2024.

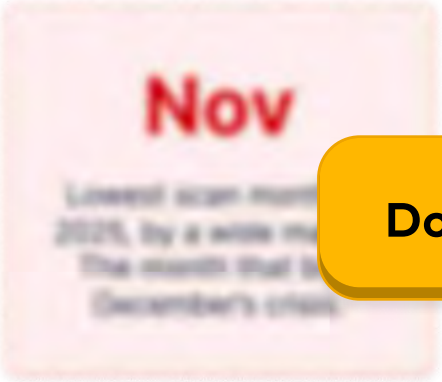

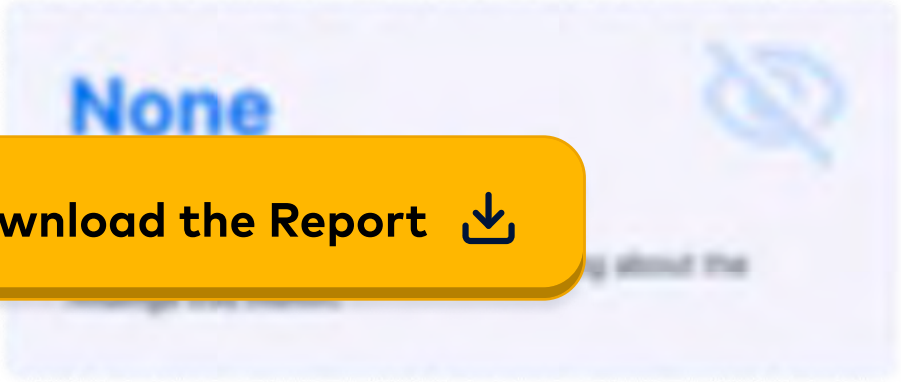

Download the Report

The moment any security leader, CISO, CTO, or even our own C-suite sees the December numbers, there is only one question: What happened in December?

However, data suggests we are asking the wrong question; everything that mattered happened in November.

This section is about a 30-day gap that sits inside every security dashboard, a structural blind spot that makes teams think they are reading the present when they are actually reading the past. Once you see it, the December surge stops being a surprise and starts being a predictable consequence of a decision made a month earlier.

The Dashboard that Reads Last Month's Risk

Ask most security teams how they know their scanning programme is working, and they will point to the same two columns: how many scans were run this month and how many vulnerabilities were found. The assumption is that these two numbers are related: scanning harder yields more findings, and a month with strong scan volume is a month with strong security coverage.

In 2025, that assumed relationship did NOT exist. The months with the highest scan volumes were far from the ones with the highest scan findings. The months with the lowest scans were not the months with the lowest findings. Looking at these two numbers side by side to read security health is like looking at the right data from the wrong number.

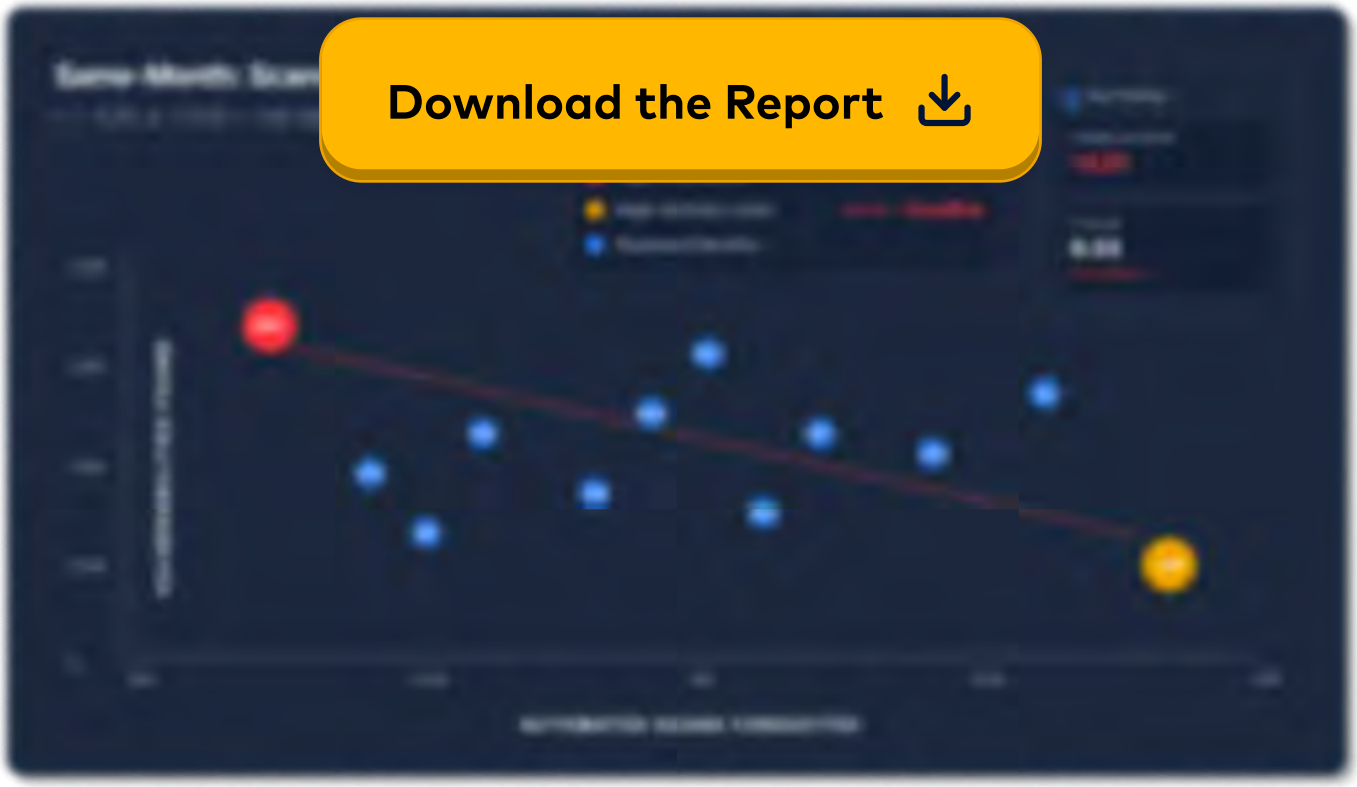

Fig 2: Every month of 2025: scans run vs. vulnerabilities found that same month. No consistent pattern exists. This is what most security dashboards display as a health signal.