HADRIAN

The verification crisis

EXTRACT OF HADRIAN'S OFFENSIVE SECURITY BENCHMARK REPORT

INTRODUCTION

More visibility was supposed to make security decisions easier. Instead, it has created a crisis of confidence. Security teams are flooded with findings, yet struggle to answer the one question that matters most: which exposures are actually exploitable right now. Our data shows that 99.5% of vulnerability scanner findings are noise and don't represent real-world risk, consuming security time, attention, and remediation capacity. The gap between visibility and verification is no longer an operational nuisance, it is a security risk.

This report extract focuses exclusively on that verification crisis. It is drawn from the Hadrian's Offensive Security Benchmark Report, which examines how modern attacks unfold and why teams must shift from volume to proof to regain control.

A visibility paradox

Security teams have significantly expanded visibility across their environments, yet confidence in decision making has declined. In our study, 95 percent of security leaders report dissatisfaction with their ability to prioritize remediation based on real world exposure. Over 70 percent said they struggle to determine which exposures are actually exploitable, and nearly two thirds cite unverified vulnerabilities as their most frustrating operational challenge.

At the same time, satisfaction with asset discovery and exposure identification is relatively high. The gap emerges after discovery. As the volume of findings increases, teams lack reliable ways to determine which issues warrant immediate action, creating hesitation rather than faster response.

TOP CHALLENGES SECURITY LEADERS FIND FRUSTRATING WHEN TRYING TO IDENTIFY AND REMEDIATE EXPOSURES

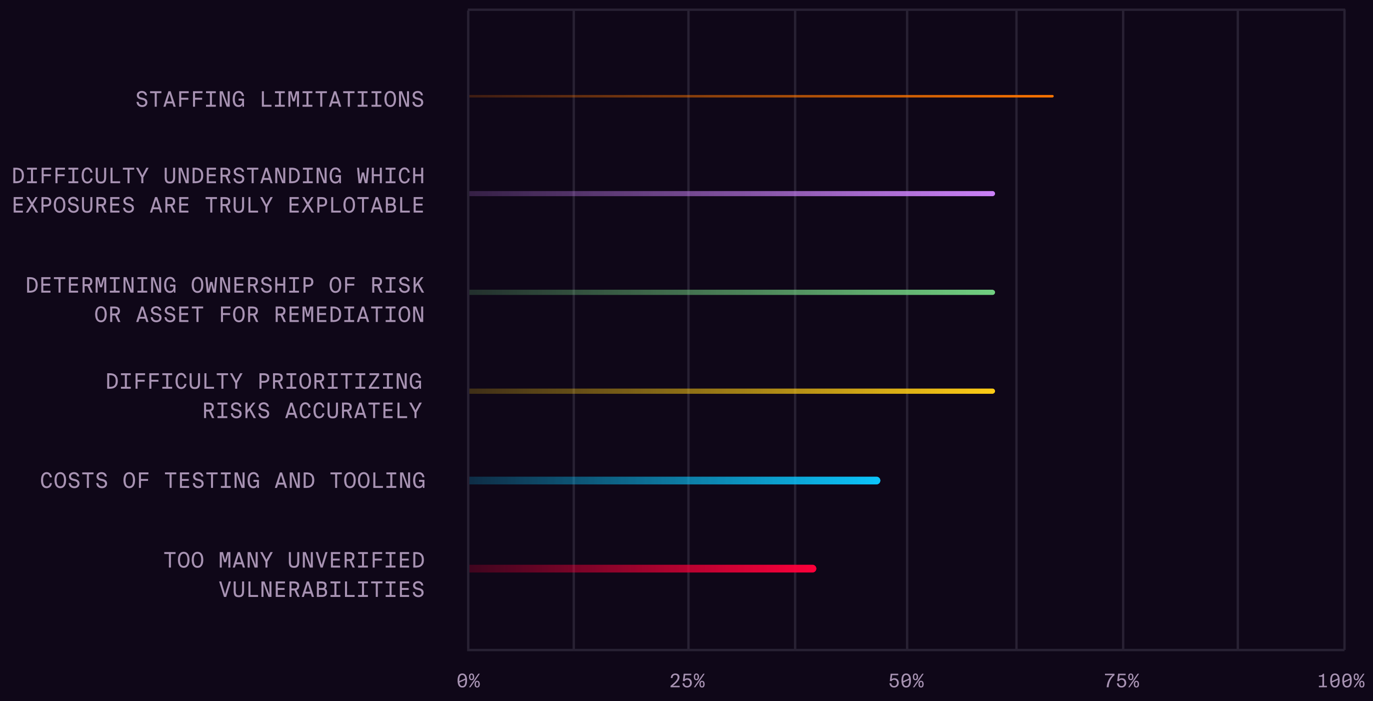

The noise problem

Security teams are not overwhelmed because they miss exposures. They are overwhelmed because almost none of what they see actually matters. On average, only 0.47 percent of risks detected by vulnerability scanners turn out to be real. Teams spend their time triaging thousands of issues that appear urgent but lack real-world impact. What remains is a flood of findings that look serious on paper but collapse under scrutiny.

Nearly 90 percent of verified risks are classified as medium or low severity. High severity issues account for just 7 percent, and critical risks make up only 3 percent of total findings. As a result, teams face thousands of findings that appear significant but lack context. The volume of medium and low severity findings obscures the small set of exposures that pose immediate, real world danger.

SEVERITY DISTRIBUTION OF VERIFIED EXPOSURES DISCOVERED IN THE ATTACK SURFACE

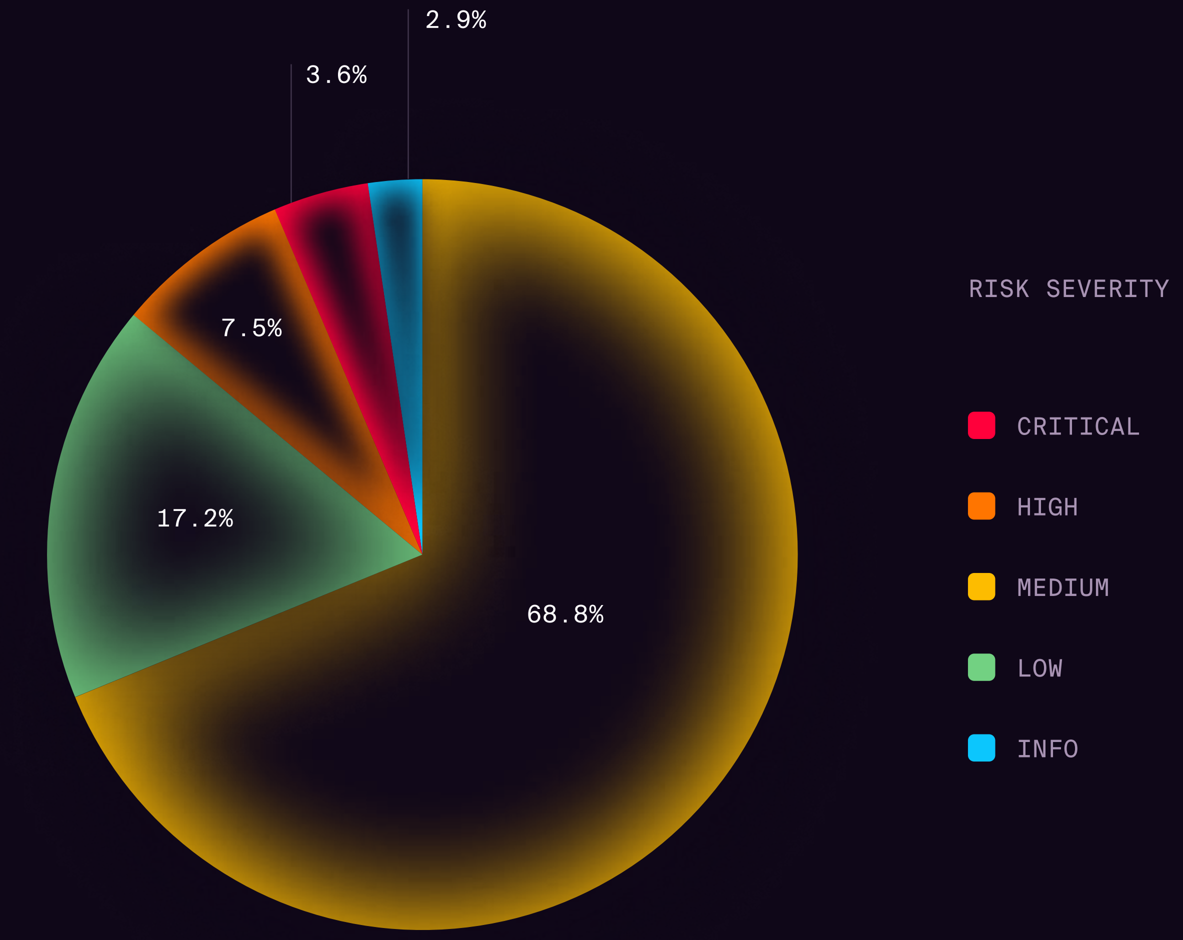

Remediation paralysis

Remediation data shows that speed is possible when the risk is clear. Critical exposures have a median remediation time of just four days. However, the average stretches to 64 days, and the slowest 10 percent take more than 120 days to resolve.

As observed last year, high severity risks take longer to remediate than less severe issues, with an average remediation time of 139 days. This is not because they matter less, but because they often require significantly more effort to fix than Medium or Low severity findings. At the same time, they lack the organizational urgency that a Critical rating creates, leaving teams caught between complexity and insufficient mandate.

AVERAGE REMEDIATION TIME BY SEVERITY

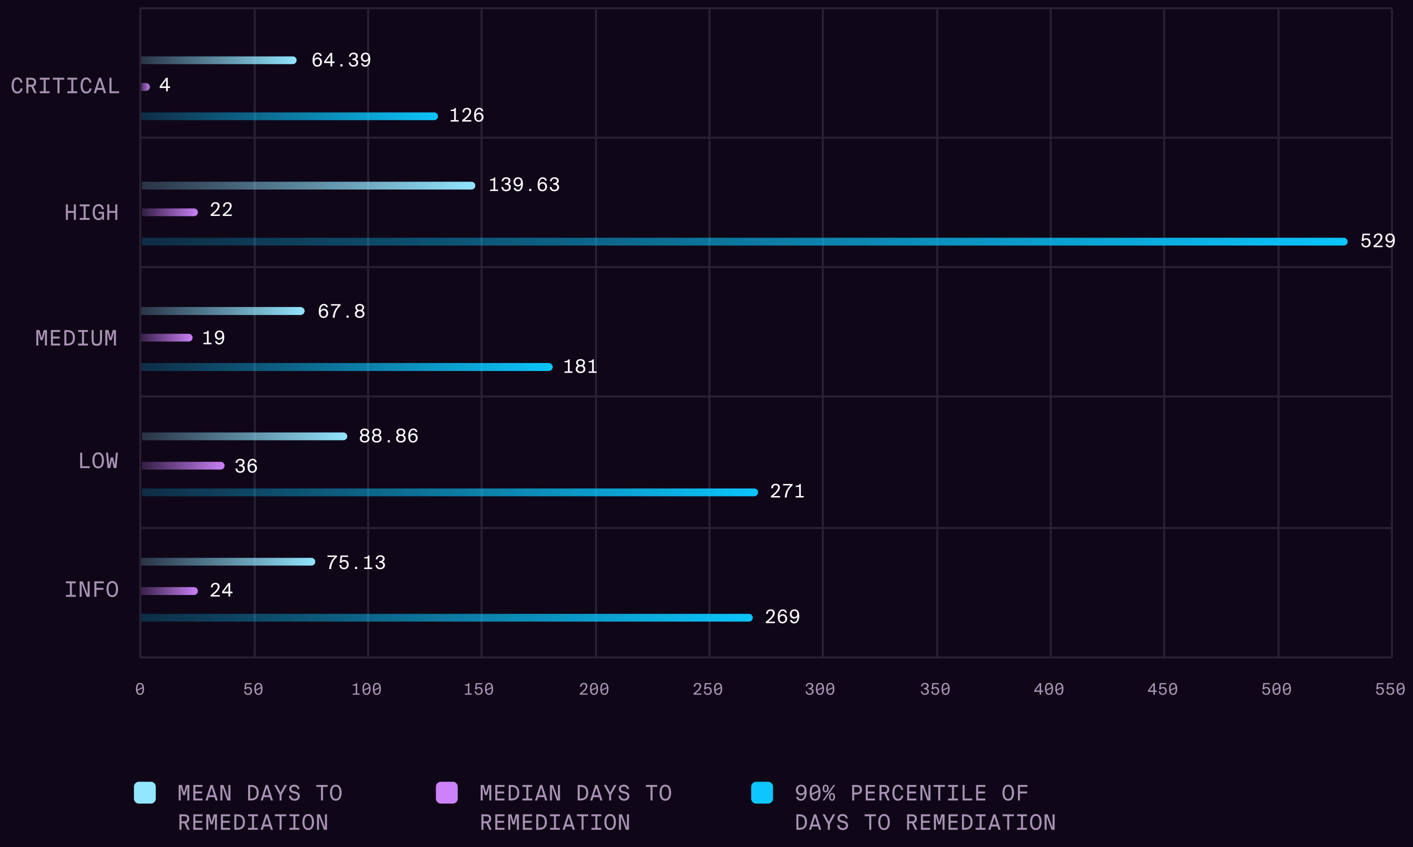

RECOMMENDATION

Security programs should shift from visibility-first approaches to verification-first security. The findings show that most detected vulnerabilities are not exploitable, yet they consume the majority of remediation effort. Teams should prioritize validating exploitability in real-world conditions, rather than flagging the presence of theoretical weaknesses in development environments. This requires integrating attack path analysis, exploitation testing, and adversary-context signals directly into exposure management workflows.

This extract covers only one part of the problem. The full Offensive Security Benchmark Report breaks down how security programs are performing and what verification-first security looks like in practice.

Get the full report

ABOUT HADRIAN

Hadrian is an offensive security platform built for teams that need proof, not predictions. We reveal exactly how an adversary could break in today by executing real attacker techniques against your external attack surface. By validating exploitability in production-safe ways, Hadrian eliminates false urgency, sharpens prioritization, and gives teams confidence to act.

Deployed in minutes and engineered for scale, Hadrian seamlessly fits into existing security workflows. Security leaders use Hadrian to move from periodic testing to continuous, evidence-backed exposure reduction.

RECOGNISED BY LEADING ANALYSTS

Gartner
GIGAOM
FROST & SULLIVAN

TRUSTED BY MARKET LEADERS

NBC
amadeus
MCKESSON
BLINQX
CRÉDIT AGRICOLE
SHV ENERGY
ABN·AMRO
London Business School
RITUALS...
SIEMENS energy
LOTTOMatica
LEROY MERLIN
WeatherTech
BIOLANDES
BLINQX
DAMEN
=exact
nedap