State of Pentesting

Report 2026

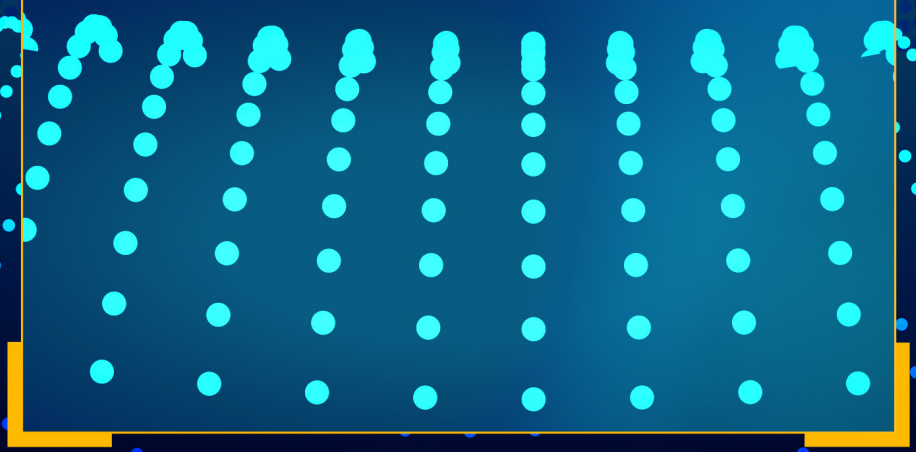

Cobalt

Thank you for reading the **2026 State of Pentesting Report**. Now in its eighth year, this report has become the premier research publication for the offensive security community, offering a depth of insight that only Cobalt—and our more than 10 years of penetration testing data—can provide.

This year's analysis draws from a massive dataset of **thousands of pentests** from a broad base of customers across industries, and a qualitative survey of **450 security leaders and practitioners**. The intersection of these two data sources helps us understand not just what vulnerabilities exist, but why they persist.

What stands out in the 2026 report from our previous research are the performance distinctions that separate the leading security teams from everyone else.

While there is universal consensus that pentesting is essential to security, the divide between organizations that handle pentest findings well and those struggling is stark. The half-life of high-risk findings—our preferred measure of remediation performance, because it accounts for both speed and completeness—ranges from **10 days** for top performers to **249 days** for the bottom tier, for **eight extra months of risk exposure**.

That **25x remediation gap** doesn't come down to resources or industry. It comes down to whether an organization treats pentesting as a one-time deliverable, or as the foundation of a continuous offensive security program.

Additionally, as we have all witnessed over the past year, an AI earthquake is shaking up broad swaths of the economy, and security is becoming collateral damage. AI and LLM applications are harboring high-risk findings at nearly **2.7 times the rate** of traditional software, yet they suffer the lowest resolution rates. This year, we explore the dilemma security teams face in enabling this game-changing technology within their organizations, while fending off AI-powered attacks.

At Cobalt, we believe that security confidence must be grounded in evidence rather than assumptions. Our **human-led, AI-powered™ approach** to pentesting as a service (PTaaS) is designed to bridge the divide between strategy and execution, transforming pentesting from a static report into a continuous risk-reduction program.

We invite you to dive into these pages to see how your organization measures up. These insights are intended to provide actionable clarity for practitioners and strategic value for executives and boards alike—helping the community move toward a safer, more resilient future.

Gunter Ollmann
Chief Technology Officer, Cobalt

Joe Brinkley
Director of Offensive Security Research and Community, Cobalt

Contents

Executive Summary

Every organization running a security program has a theory about how exposed they are to risk. This report is about what the data actually shows.

The 2026 State of Pentesting Report draws on two sources that few research efforts can claim simultaneously: a survey of 450 information security leaders and practitioners, and the results of thousands of penetration tests conducted over five years. The survey tells us what security teams believe, intend, and prioritize. The pentest data tells us what they find and what they do about it. Together, they reveal a richer, more honest picture of the security landscape than either could alone.

That pairing produces some of the report's most instructive moments. Security leaders feel increasingly good about their programs: budgets are growing, pentesting is universally viewed as essential, and more organizations are taking a programmatic approach to offensive security. Yet the pentest data reveals persistent gaps: high-risk findings that go unresolved for months, remediation timelines that consistently exceed the service level agreements (SLAs) organizations set for themselves, and serious challenges to remediate vulnerabilities in AI applications in spite of accelerating adoption.

The report that follows is organized to take you from the big picture to the specific:

- Practitioner-surveyed context for the state of offensive security programs.
- Demonstrating what AI and LLM application pentesting is revealing about an attack surface that is growing at a pace security practices can't match.
- Industry benchmarks for how organizations measure up on the metrics that matter.
- The most consequential vulnerabilities our pentesters are finding across web, mobile, and network environments.
- Our Offensive Security Leaders Quadrant, defining benchmarks at the intersection of program strategy and SLA execution.
- Recommendations for maturing your program.

The data is here to help you on your journey. **Let's get into it.**

Key Findings

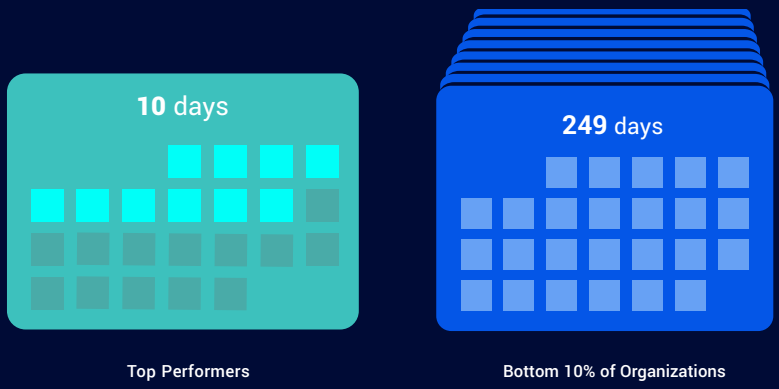

8 months

Increased vulnerability exposure time for organizations in the bottom 10%

While top-performing organizations achieve a high-risk finding half-life—the key metric for resolving critical issues—of just **10 days,** high-risk vulnerabilities in the bottom tier languish for **249 days**, a difference of eight months.

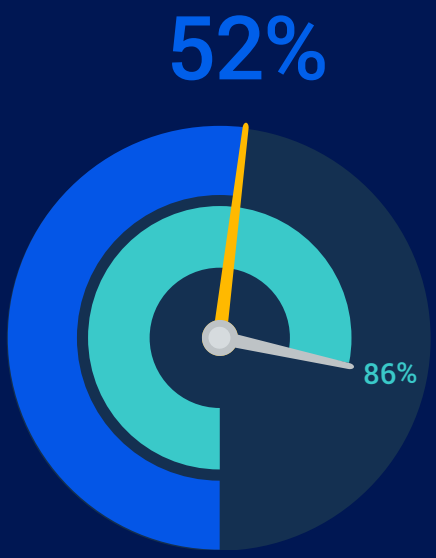

Overall resolution rate of high-risk findings

While the typical organization resolves **86%** of its high-risk findings, the total resolution rate across the entire five-year dataset is stuck at just **52%.** This suggests some teams are cherry-picking easy fixes, while leaving nearly half of their long-term foundational risks open to exploitation.

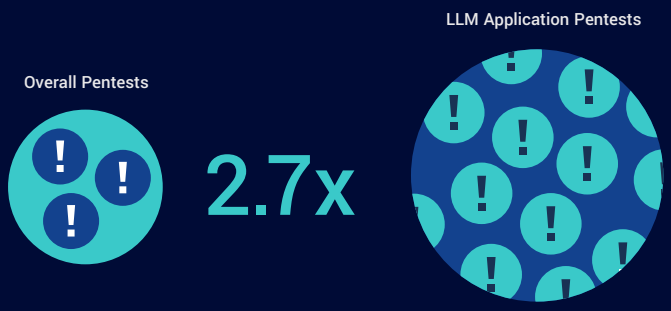

Increased rate of high-risk vulnerabilities in LLM application pentests vs. pentests overall

Our pentest data reveals that **32%** of all AI/LLM findings are rated as **high risk**—nearly **2.7x the rate** found in the overall dataset **(12%)**. Not surprisingly, security teams' confidence in AI security plummeted since last year—dropping from **64%** in 2025 to **51%** this year.

Resolution rate of high-risk vulnerabilities in AI/LLM tests

AI/LLM pentests show a resolution rate of only **38%**, the **lowest of any testing category**. Although this is an improvement from last year's findings, it's a security loophole that attacks like prompt injection are liable to exploit.

Advantage of programmatic approach in achieving 3-day SLAs vs. compliance-only testing

Organizations that adopt a programmatic approach to offensive security are **4.5x** more likely to resolve critical findings in **three days or less** compared to those operating under a compliance-driven or ad hoc model.

15% Security practitioners who believe their teams consistently meet critical SLAs

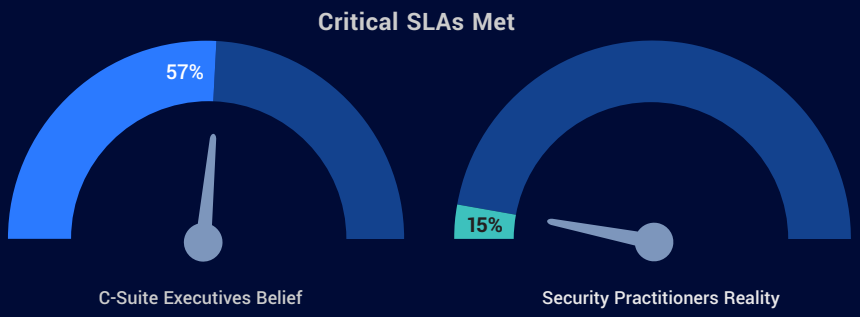

57% of C-suite executives in our survey believe their organization consistently meets remediation SLAs, yet only **15%** of security practitioners who actually perform the work agree. This disconnect underscores the reality of the engineering bottlenecks and resource constraints that stall remediation.

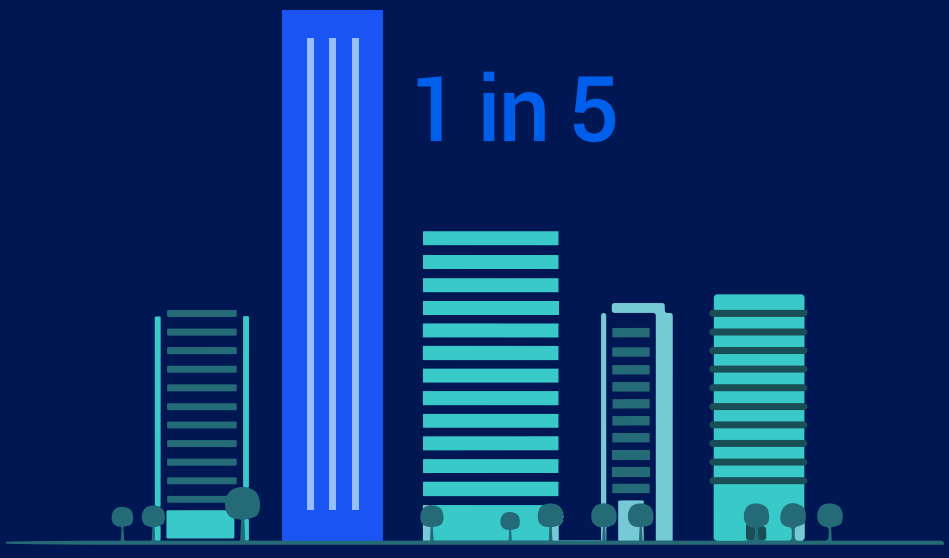

Organizations have had an AI-related security incident

19% of organizations admit to AI-related security incidents that can result in data breaches. Although **44%** said they have never suffered an AI security incident, a further **18%** are unsure, and **19%** prefer not to answer.

Pentesting Perspectives

The security industry generates no shortage of opinions about what organizations should be doing. Survey data offers something rarer: a window into what executives and practitioners on internal security teams actually believe—and how they're actually behaving.

We surveyed 450 information security leaders (C-suite and VPs) and managers/practitioners, and what they told us paints a picture of a discipline that is maturing, expanding its mandate, and grappling with a threat landscape that keeps raising the stakes.

Pentesting is essential to modern security

Start with the baseline: 97% of respondents view pentesting as foundational to modern security programs, up three points from last year. That near-unanimity is striking on its own, but the budget data makes it concrete.

Eight in 10 organizations report that their offensive security budgets grew in the past year—nearly a third say significantly and half say incrementally. Security leadership isn't just saying pentesting matters; they're funding it.

Pentesting is a business enabler

Pentesting has become a commercial signal, not just a compliance-driven security control. Three-quarters of respondents say pentesting improves customer trust in their products, and the same proportion say it improves product quality and innovation.

Perhaps most telling: **61% report that customers are now actively requesting third-party pentest reports** to validate software security, up 13 points from last year, ranking second only to compliance certifications as a proof mechanism for security posture.

Pentesting is becoming more programmatic

The shift in motivation for pentesting is reshaping how organizations structure their programs. For the first time in this survey series, the share of organizations taking a programmatic approach to pentesting (continuous, integrated, risk-driven) exceeds the share pentesting primarily to satisfy compliance requirements: **53% versus 40%.**

There is a caveat: executives are significantly more likely than practitioners to describe their approach as programmatic (63% vs. 42%). Whether that gap reflects genuine organizational reality or aspirational framing from the top, executive belief tends to drive resourcing decisions, and resourcing decisions eventually shape practice.

GenAI adoption is driving new demand

Fifty-eight percent of organizations say they have an explicit strategy for genAI adoption, though most are still mid-rollout, and just **8% describe their genAI strategy as fully adopted.**

As organizations embed AI capabilities into products and infrastructure, the attack surface evolves in ways that standard testing approaches weren't designed to probe. Cobalt conducted 2.4 times more AI/LLM pentests in 2025 than in 2024. That acceleration reflects demand. What it's finding, and what it means for security programs, is the subject of the next section.

Pentesting AI

Artificial intelligence has moved from boardroom buzzword to operational reality faster than most security programs can keep pace. Organizations are embedding genAI into products, processes, and pipelines at a sprint.

The security function is often left jogging to catch up to the possibility that these AI applications become a Trojan horse—ostensibly providing value to the company, but with hidden vulnerabilities lurking inside. Not to mention that adversaries are using AI themselves to accelerate and automate attacks.

That dynamic is clearly visible in this year's data. Survey respondents and pentest results tell a consistent story: AI adoption is accelerating, AI-related threats are intensifying, and the security practices designed to counter those threats are lagging behind.

AI is transforming the threat landscape

A striking 93% of surveyed cybersecurity professionals report observing an increase in genAI-related threats from external actors—up 12 percentage points from last year. And these aren't just more polished phishing lures that end up in the spam folder. The threats are translating into actual incidents.

Has your organization had an AI-related security incident?

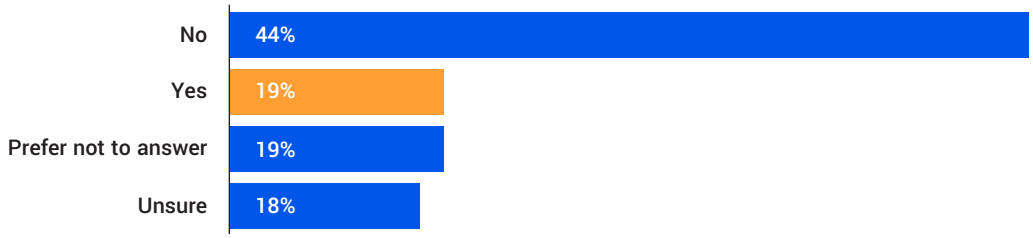

Figure 1

Source: Cobalt State of Pentesting Report 2026

Nearly one in five organizations (19%) has experienced an AI- or LLM-related security incident. And that figure almost certainly understates reality: 18% of respondents said they weren't sure whether they'd had an incident, and another 19% declined to answer the question. When a third of your survey population can neither confirm nor deny AI-related incidents, the true exposure is likely higher than the headline number suggests.

What was the nature/cause of your AI-related incident?

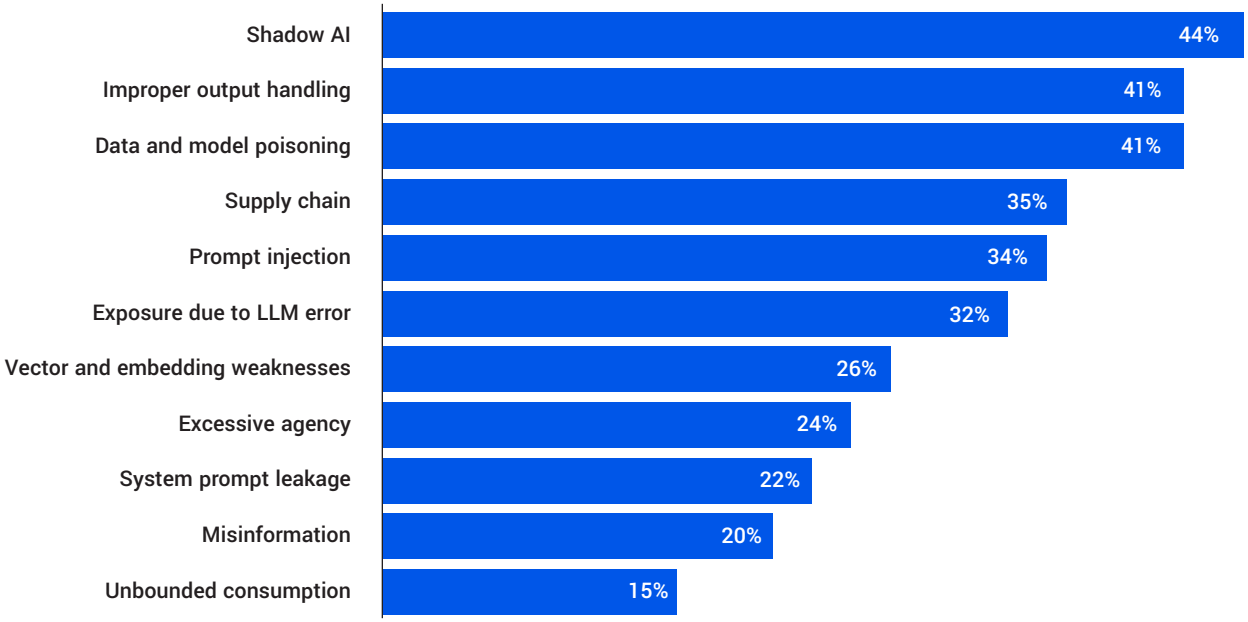

Figure 2

Source: Cobalt State of Pentesting Report 2026

When incidents do occur, the causes span a range of vectors. Shadow AI tops the list—unauthorized AI tools adopted outside any security review process—cited by 44% of those reporting incidents. Data and model poisoning (41%) and improper output handling (41%) follow closely, with software supply chain vulnerabilities (35%) and prompt injection (34%) rounding out the top five. The breadth of this list matters: there is no single chokepoint to defend. AI risk is distributed.

Penetrating Insight

19% of organizations have experienced an AI/LLM-related security incident, and roughly a third of respondents couldn't or wouldn't confirm whether they had. The true prevalence of AI-related breaches is almost certainly higher than self-reported figures indicate.

Confidence is dropping as complexity rises

Here is where the data takes a meaningful turn from last year. Security practitioners aren't simply alarmed about AI threats in the abstract. Their confidence in their own ability to handle those threats is declining.

About half of organizations believe they are well-equipped to address the security implications of AI adoption. That sounds reasonable until you notice that figure is down 13 percentage points from last year. At the same time, 61% say there is a need for a strategic pause to recalibrate and reinforce defenses against AI-driven threats. That's up 13 percentage points year-over-year.

It's not just a coincidence that confidence declined by the same margin as the calls for a timeout increased. This reflects a growing recognition that the gap between AI adoption and AI security is widening, not closing. Organizations are increasingly aware of what they don't know—which is, at least, a healthier posture than false confidence.

That self-awareness surfaces in another finding: 62% of respondents say they need improved capabilities to properly test the security of genAI tools. That's the right instinct. The pentest data that's coming up soon shows exactly why.

What are your views of genAI in the context of cybersecurity?

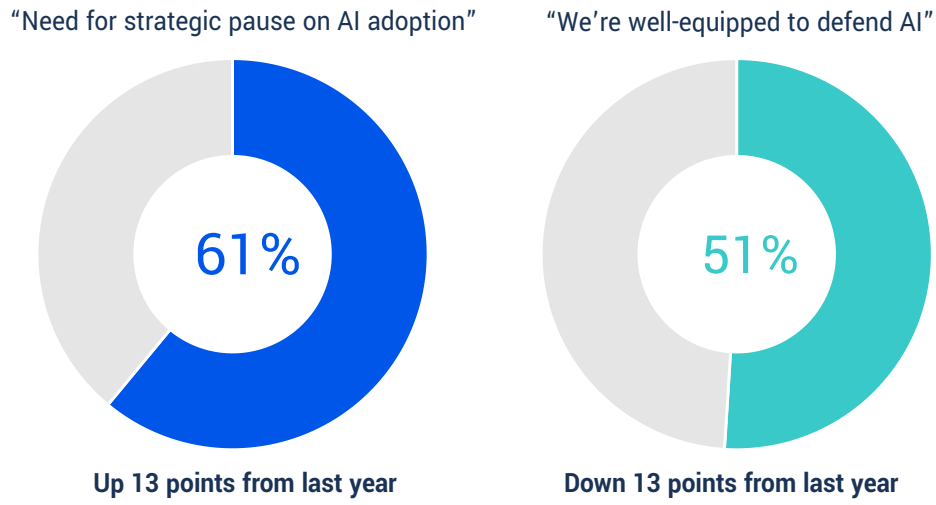

Figure 3

Source: Cobalt State of Pentesting Report 2026

Security teams are rising to the AI challenge

Despite the anxiety, organizations aren't standing still. The survey shows meaningful strides in both defensive monitoring and offensive security testing aimed at AI systems.

- 83% of firms now monitor AI behavior for malicious use or abuse—up 21 percentage points from last year.
- 77% conduct regular security assessments and pentests on genAI products—up 11 percentage points.
- Over half use red teaming or adversarial testing focused specifically on LLM security.

These are encouraging numbers, but they need to be read alongside the incident data: despite wider AI behavior monitoring, nearly one in five organizations still experienced an AI-related incident last year.

Monitoring is necessary but not sufficient. And the gap between organizations conducting assessments (77%) and those confident in their AI security posture (~50%) signals that testing is surfacing problems faster than organizations can remediate them.

What's your organization doing to secure genAI products?

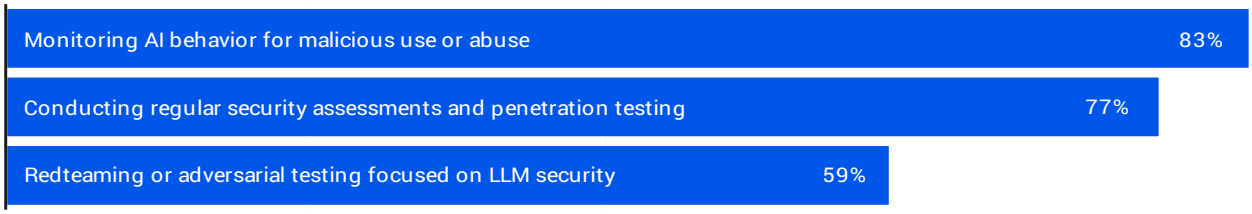

Figure 4

Source: Cobalt State of Pentesting Report 2026

What pentesters find in AI and LLM applications

Survey data captures perceptions and reported practices. Pentest data captures what's actually there. The two perspectives are illuminating in combination, and in this case, they tell a coherent story about an attack surface that is still maturing.

Prompt injection vulnerabilities are by far the most common finding in AI pentests—and that dominance only widened over the last year. Safeguards are being added to and built around AI applications, and these vulnerabilities allow adversaries to manipulate model inputs to bypass them. Given the AI adoption rate, that's unlikely to abate anytime soon.

Relative frequency of findings from AI/LLM pentests

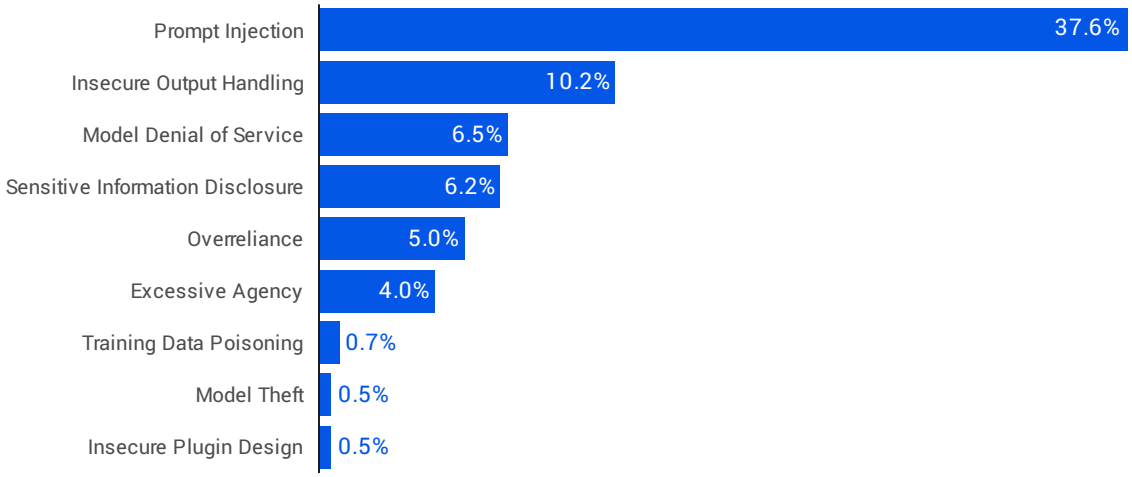

Figure 5

Source: Cobalt State of Pentesting Report 2026

Insecure output handling is next on the list. When models generate responses that are passed downstream without proper validation, the results can include data leakage, injection attacks, or security misconfigurations. Model denial-of-service (DoS) and sensitive information disclosure are neck-and-neck among AI-specific findings. Over reliance captures a subtler risk: the tendency for incorrectly generated or hallucinated content to create misinformation, legal exposure, or important decisions made on faulty premises.

OWASP recognized the evolving nature of these risks and updated its Top 10 for LLM and GenAI for 2025. Notably, the DoS category was expanded to 'Unbounded Consumption,' which incorporates Denial of Wallet (DoW) attacks (exploiting the cost-per-use model of AI services to drive up compute costs as an attack vector). This is an emerging risk category with no real analog in traditional application security, and one worth watching as AI API usage scales.

AI findings are disproportionately serious

Not all pentest findings carry the same weight. Throughout this report, we focus particular attention on high-risk findings. We'll get into severity ratings later, but for now, suffice it to say that high-risk findings are the ones most likely to result in harm to the business. Across all pentest types in our data-set, roughly 12% of findings meet that threshold.[1]

For AI and LLM pentests, that proportion is more than double: approximately one in three findings (32%) is rated high risk. The reasons for this are structural. Vulnerabilities in AI applications aren't a problem that enabling automated patching can make go away. The attack surfaces are novel, the failure modes are not fully characterized, and many development teams building on top of LLMs lack the security training to anticipate what adversaries will try.

What should organizations take from this elevated rate of high-risk findings? The intuitive response—avoid AI—is neither practical nor useful given the pace of adoption. The more productive framing: if you're deploying AI, you are almost certainly increasing exposure unless you're actively finding and resolving vulnerabilities. Unfortunately, many are struggling in the resolution department.

[1] See Figure 8 in the next section for a breakdown of prevalence/severity ratings among findings.

Elevated prevalence in high-risk findings in AI pentests

Figure 6

Source: Cobalt State of Pentesting Report 2026

Penetrating Insight

AI applications inherit all the vulnerabilities of traditional applications, then add a new layer of LLM-specific risks on top. Security teams can't treat AI pentesting as a replacement for standard application testing—it's additive.

The resolution gap: most AI findings go unfixed

High severity is only half the problem. The other half is what happens (or doesn't happen) after findings are discovered.

You might expect that the high severity of AI findings would drive correspondingly high remediation urgency. The data shows the opposite. AI and LLM pentests have the lowest resolution rate of any pentest method Cobalt conducts.

Resolution rate of high-risk findings by pentest type

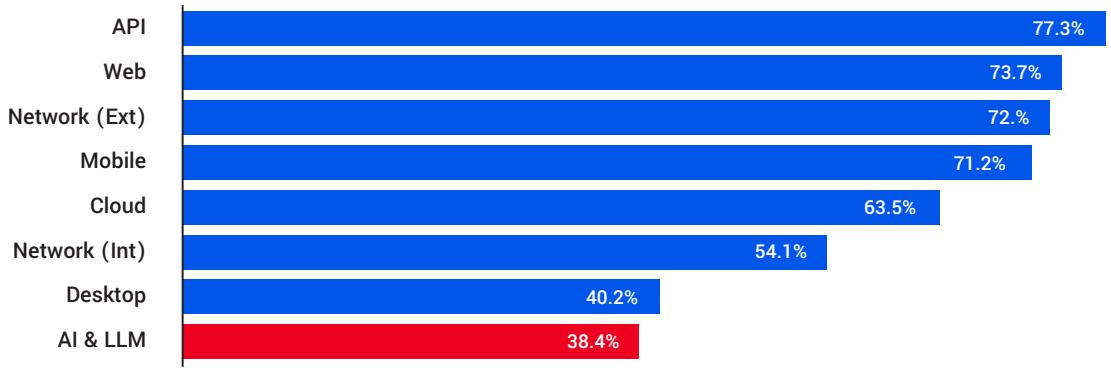

Figure 7

Source: Cobalt State of Pentesting Report 2026

On the bright side, the resolution rate for high-risk AI/LLM findings has risen from 21% to 38% over the last year. That's real progress and worth acknowledging. Even at 38%, though, the picture is stark. For every risky AI vulnerability resolved, two remain open and exploitable.

Three factors help explain this gap, and they're unlikely to disappear anytime soon:

- Knowledge deficit. Fixing AI/LLM vulnerabilities requires a combination of security expertise and AI/ML expertise that few teams currently have in-house.
- Vendor dependency. Some AI findings can't be fixed by the organization that discovered them. When the vulnerability lies in the underlying model rather than in the application built on top of it, the path of resolution runs through the LLM vendor.
- Organizational novelty. AI applications are often newer initiatives with less mature security processes, less integration with vulnerability management systems, and less organizational muscle memory around what to do when a pentest surfaces a finding.

Penetrating Insight

AI and LLM pentests produce high rates of risky findings and the lowest rate of resolution of any pentest type. That gap is closing, but at the current trajectory, it will remain a wide risk exposure for years to come.

Benchmarking Performance

The pentest reports we deliver to clients revolve around the security issues uncovered, which makes perfect sense. They've contracted the pentest to discover and fix unknown vulnerabilities, so they should be upfront and center. For this annual "State of" report, however, we're going to mix things up and move those details to the end.

Here we step back from the specifics of any single test to examine high-level performance patterns across thousands of organizations. How risky are the findings uncovered? How effective are organizations at resolving them? How quickly does that resolution happen? What are the performance metrics of leading vs. lagging organizations?

The vulnerability details come later. This section establishes key benchmarks you can use to mature your offensive security program.

As we review these benchmark stats, it's worth noting that pentests typically target specific applications or scopes. Unlike traditional vulnerability scans, which can uncover tens of thousands of vulnerabilities across assets in a large environment, pentests put more focus on enumerating all the ways an attacker would actually get in rather than theoretical weaknesses.

We're examining real risk to the business rather than the noise of scanners.

Severity of findings

Not all security issues discovered during a pentest carry equal weight. Cobalt rates each finding on two dimensions: the likelihood of exploitation and the potential impact on technical and business operations. The combination of those two ratings determines severity, from low-level informational observations to critical security exposures.

A proportional breakdown of severity for all findings in our sample is portrayed in Figure 8.[2] The upper-right region of that matrix—findings with both high likelihood and high impact—represents the exposures that warrant fast action and will receive the most attention throughout this report. We call these high-risk findings.

High-risk findings account for roughly 13% of all non-informational pentest findings, and that's been remarkably stable over time. That's good news for organizations trying to focus on the subset of findings that pose the greatest risk. It's not a rapidly-moving target, which makes the prevalence of high-risk vulnerabilities a reliable baseline against which to measure performance.

[2] Informational findings (rating of very low on both scales) have been removed from the chart.

Severity rating of pentest findings and breakdown by severity

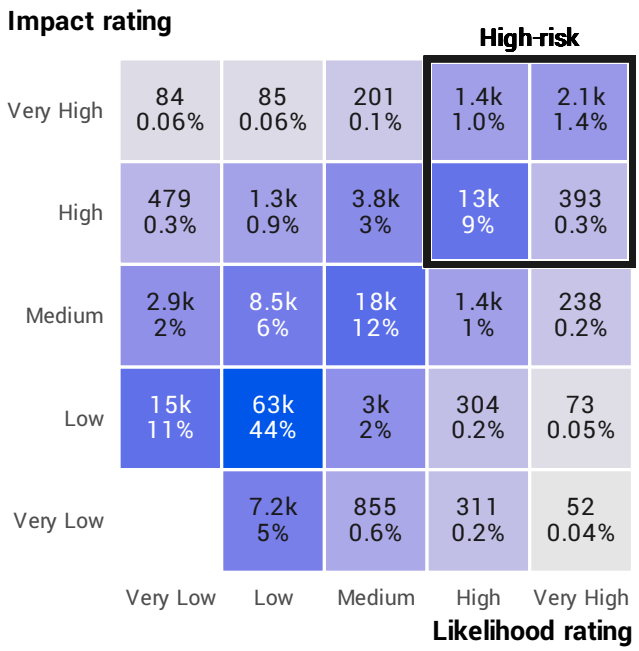

Figure 8

Source: Cobalt State of Pentesting Report 2026

Prevalence of high-risk findings among top-10%, bottom-10%, and median organizations

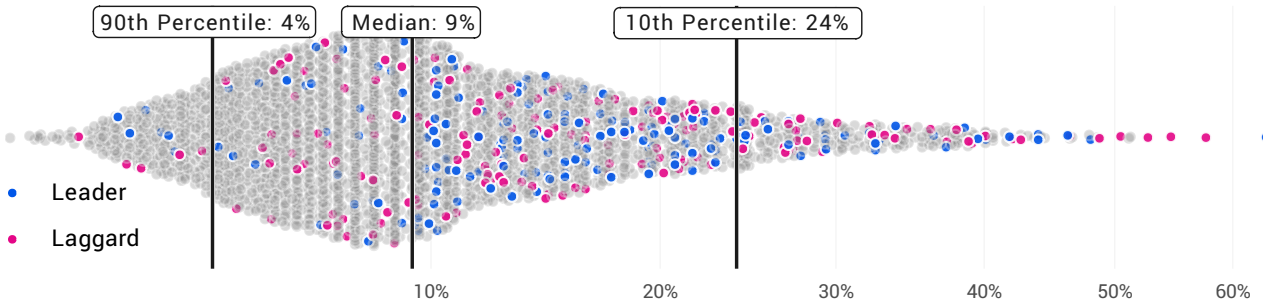

Figure 9

Source: Cobalt State of Pentesting Report 2026

The previous chart examines the proportion of all findings assessed as high risk. But if we want to establish a baseline for assessing risk and performance, we need to measure this at the organizational level. Figure 9 presents this view, with each dot representing an organization and plotted along the x-axis according to the proportion of its findings rated high risk.

The median ratio of high-risk findings is 9%. Most organizations cluster within roughly 10 percentage points of that median in either direction. But the distribution has a long right tail: a subset of organizations contends with a significantly elevated rate of risky exposures.

Wondering about the significance of the colors in Figure 9? Great! The blue dots indicate high-performing organizations that are both fast and thorough at resolving high-risk findings. The red dots represent firms on the other end of that performance spectrum; they're struggling with a long backlog of risky issues. The specific metric we're using to make this performance distinction is half-life,[3] which we will define in a later section.

We wanted to distinguish top and bottom performers here because there's a very important observation we want to make: they're intermingled when it comes to the ratio of high-risk findings. That means you're not destined to fail if you start out with a lot of risky issues. Conversely, you're not guaranteed success if high-risk findings are few and far between on your pentest report.

On that note, it's worth clarifying that a high rate of risky findings doesn't necessarily signal insecure development or lax remediation. The nature of the assets tested affects this rate, and different types of pentests tend to identify issues of varying volume and severity.

We'll focus most of the forthcoming analysis on these high-risk findings. Why? Well, the survey data provides solid justification for that. When asked which factors most influence their decisions about prioritizing pentest findings, security leaders and practitioners ranked potential business impact first—by a wide margin.

[3] For convenience: Half-life is the time to resolve half of all findings based on survival analysis.

Most influential factors for resolving pentest findings

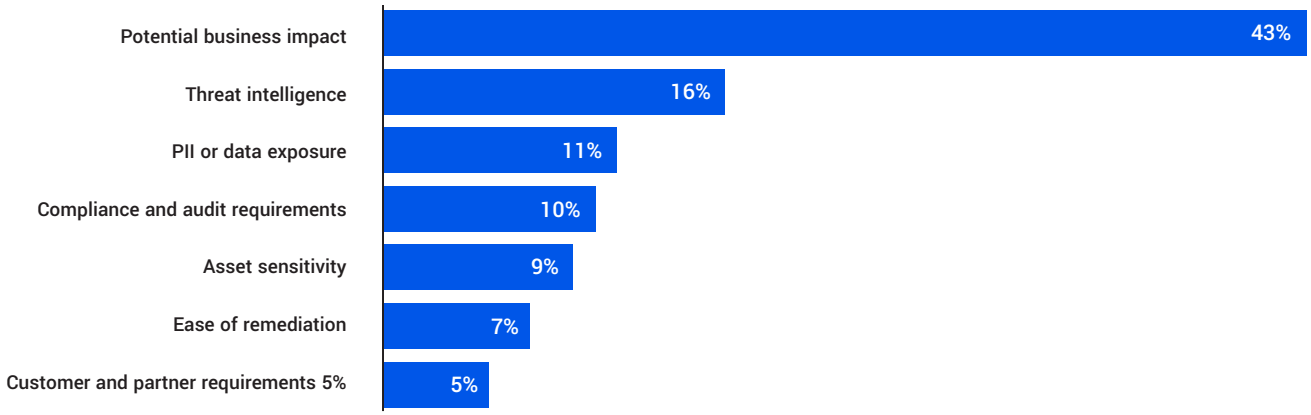

Figure 10

Source: Cobalt State of Pentesting Report 2026

It's also worth noting that the second- and third-most-cited factors, threat intelligence and exposure of PII or sensitive data, are embedded in how we define high-risk findings. Threat intelligence informs the likelihood rating; data exposure drives impact. So focusing our benchmark analysis on high-risk findings isn't an arbitrary editorial choice; it reflects how practitioners themselves think about what matters.

Penetrating Insight

Most organizations start from roughly the same place with about one in ten findings rated as high-risk, i.e. exactly the kinds of issues pentests are designed to uncover. The differences between top and bottom performers only begin to emerge when we look into remediation.

Resolution of findings

Discovering security issues is the first step, but reducing risk requires resolving those findings. That's easier said than done. Security teams are pulled in multiple directions simultaneously, and pentest findings compete for attention with alerts, audit requests, and a backlog that rarely shrinks fast enough.

Still, pentest findings tend to rise above the noise. The majority of survey respondents place pentest findings among their top three remediation priorities, alongside CISA Known Exploited Vulnerabilities (KEVs) and alerts from XDR/EDR tooling. That prioritization shows up in the numbers.

The typical organization resolves 86% of its high-risk pentest findings. Just 1% of organizations resolve less than 10%. The leading (blue) and lagging (red) performers show strong separation here.

Resolution of high-risk findings among organizations

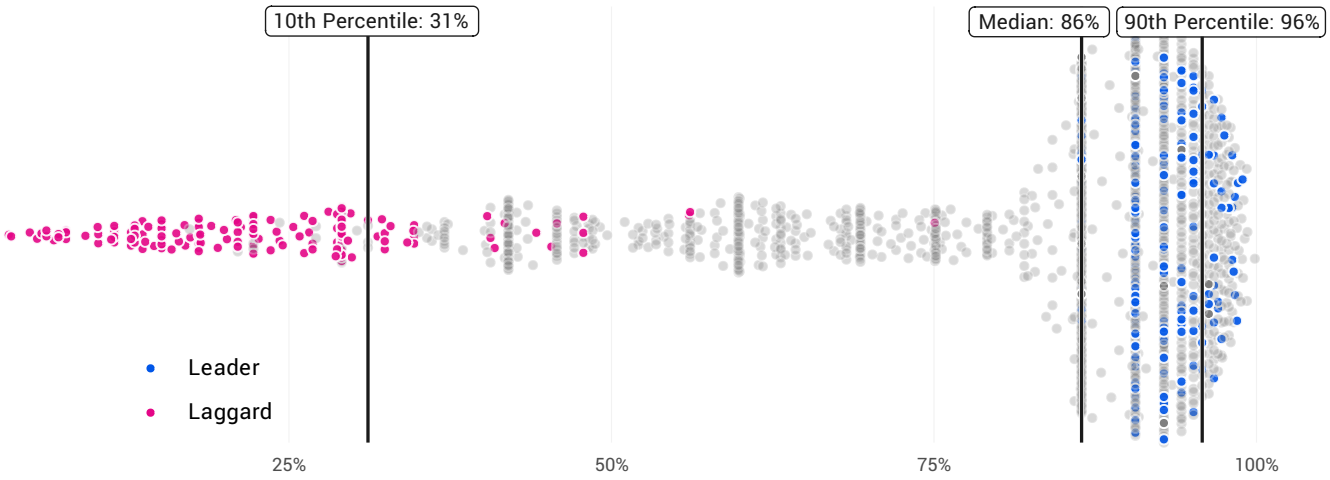

Figure 11

Source: Cobalt State of Pentesting Report 2026

If you look back at the 2025 edition of the State of Pentesting Report, you'll see a resolution rate based on the overall percentage across all pentests and findings. The comparable stat for the most recent year is 52%, which is an all-time high. We decided to focus on the organization-level view in this report to establish a benchmark and emphasize performance differences.

While the resolution rates depicted here are quite high (for most firms, at least), it's probably a conservative view of reality. They count any finding that hasn't been explicitly marked as resolved or risk-accepted in the Cobalt platform as unresolved. That includes findings Cobalt has actively confirmed as open, as well as findings where the customer simply hasn't reported back. Cobalt offers free retesting to validate fixes, but not all customers take advantage of it. For findings in that latter group, "unresolved" more accurately means "no evidence of resolution." The true fix rate is likely somewhat higher.

That said, there are legitimate reasons to leave high-risk findings open. The survey sheds light on the most common reasons in Figure 12. Concerns over business disruption were cited by 41% of respondents. Another 28% point to compensating controls that, in their assessment, neutralize the risk of exploitation. Rounding out the top 3 reasons, 18% said the cost of remediating security issues is a deterrent to addressing them.

None of these reasons should be read as excuses. They reflect the real-world constraints security teams operate under. A finding left open with a documented compensating control and a business-impact rationale is a materially different situation from a finding left open through inattention. The data doesn't distinguish between the two, which means the 19% of organizations resolving fewer than half their high-risk findings includes both deliberate risk acceptances and genuine remediation failures.

Reasons cited for not remediating high-risk findings

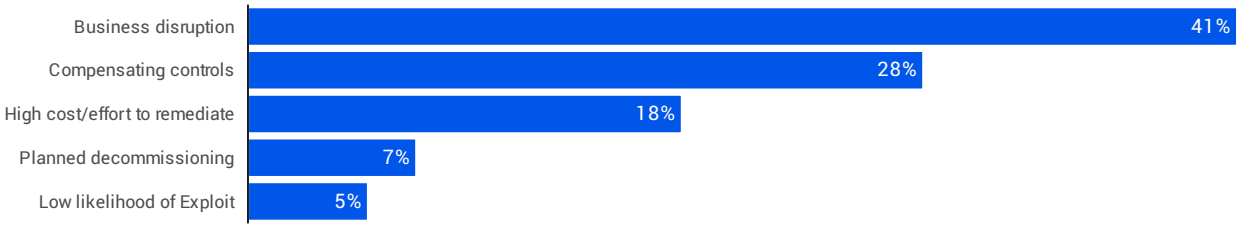

Figure 12

Source: Cobalt State of Pentesting Report 2026

Penetrating Insight

The typical organization resolves 86% of high-risk pentest findings, but that median statistic masks wide dispersion. A subset of organizations resolves very few findings, pulling down overall statistics and representing persistent risk exposure.

Mean time to resolution

We now know that most organizations resolve most of their pentest findings. But how quickly does that happen? Enter mean time to resolution (MTTR), which measures the speed at which pentest findings are resolved (at least for those that are resolved). But before measuring actual speed, it's worth examining the speed organizations aim for.

Adoption and achievement of formal SLAs for remediation

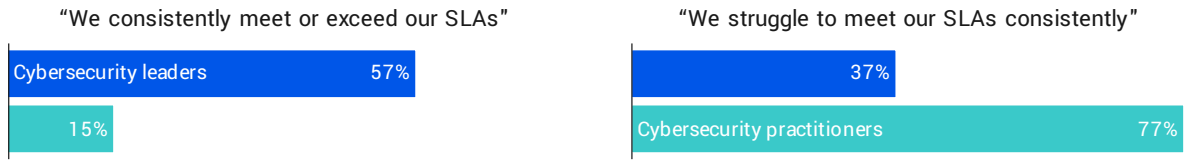

Figure 13

Source: Cobalt State of Pentesting Report 2026

What SLAs stipulate

About a third of organizations report having formal SLAs for vulnerability remediation and claim to be meeting them consistently. More than half have SLAs but acknowledge they struggle to hit them with regularity. The gap between setting a target and hitting it is one of the defining features of security operations in practice.

What makes this finding particularly interesting is who's speaking. Executives and practitioners tell starkly different stories about SLA adherence: 77% of practitioners say meeting SLAs is a genuine struggle. Among executives, that number drops to 37%.

The gap likely reflects how removed leadership can be from day-to-day operations. We suspect the view from the trenches is probably closer to reality.

The SLA targets themselves are ambitious. Half of organizations aim to remediate critical vulnerabilities within a week. Targets for high and medium vulnerabilities are somewhat less aggressive, though still short. The pentest data that follows will show how those aspirations measure up against reality. The answer is that even the most lenient SLA targets generally fall short of what organizations actually achieve.

SLA time requirements for resolving critical findings

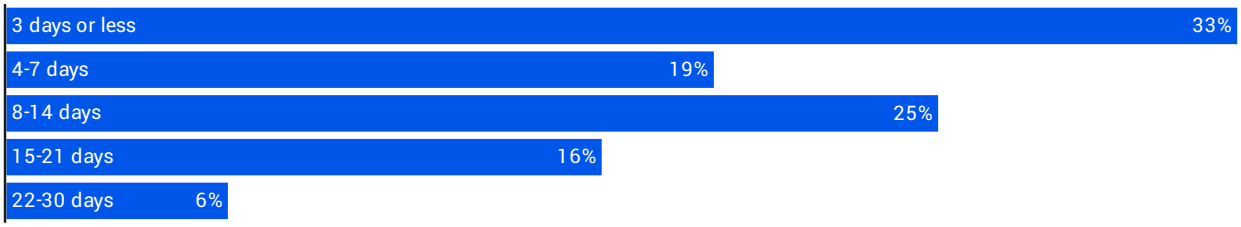

Figure 14

Source: Cobalt State of Pentesting Report 2026

What pentest data shows

Let's move on from aspirations to actual performance. The pentest data clocks the median MTTR across all organizations at 39 days to resolve high-risk findings. The fastest 10% gets it done within 13 days or fewer, while firms with the slowest MTTR stretch that timeline to 131 days or more.

If you read the 2025 report, you may recall us making a distinction between MTTR and half-life and advocating for the latter as a better overall performance measure. Figure 15 offers justification for that stance. Notice how there are some organizations that score very poorly when it comes to half-life (red dots), yet rate well above average for MTTR. These organizations address *some* issues quickly, but they leave the majority unresolved. They're sprinting but never get close to the finish line.

MTTR of high-risk findings among organizations

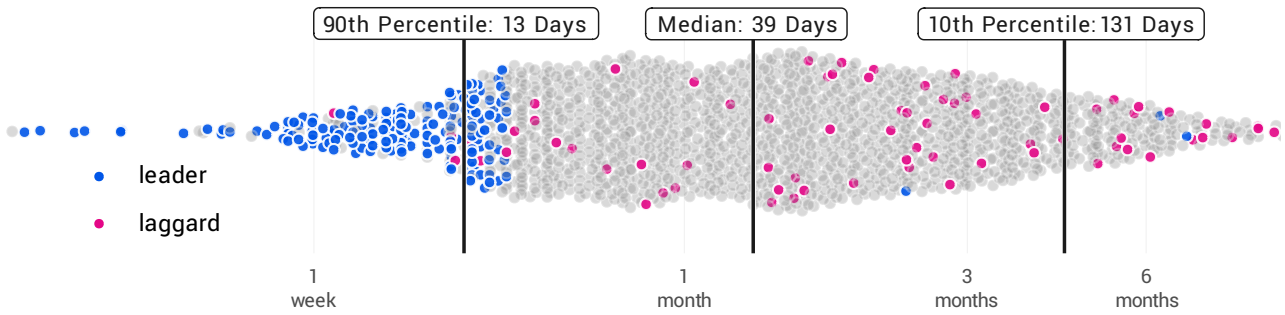

Figure 15

Source: Cobalt State of Pentesting Report 2026

The comparison to SLA targets is unambiguous: the median organization's MTTR of 39 days exceeds even the most lenient formal SLA targets. A small subset of organizations exhibit MTTRs that approach the 7-day critical vulnerability SLA that half of the survey respondents claim to target. The gap between intention and outcome is real and consistent.

But if recent trends hold, that gap may shrink. Figure 16 appears to indicate that high-risk MTTR has declined by 18 days in the last year. We put "appears" in quotes because the recent drop is due in part to 2025 findings having a shorter resolution history than prior years. This compresses the observed MTTR downward relative to earlier years in ways that are not entirely attributable to improved remediation practices. The trend is directionally encouraging but should be interpreted cautiously. The true year-over-year improvement is likely smaller than the chart implies.

Even with that caveat applied, it's clear that many organizations can achieve resolution times that keep pace with the SLAs reported earlier. In our experience, adopting programmatic pentesting approaches—moving from ad hoc, compliance-driven tests to continuous, integrated offensive security programs—creates the organizational inertia for faster remediation.

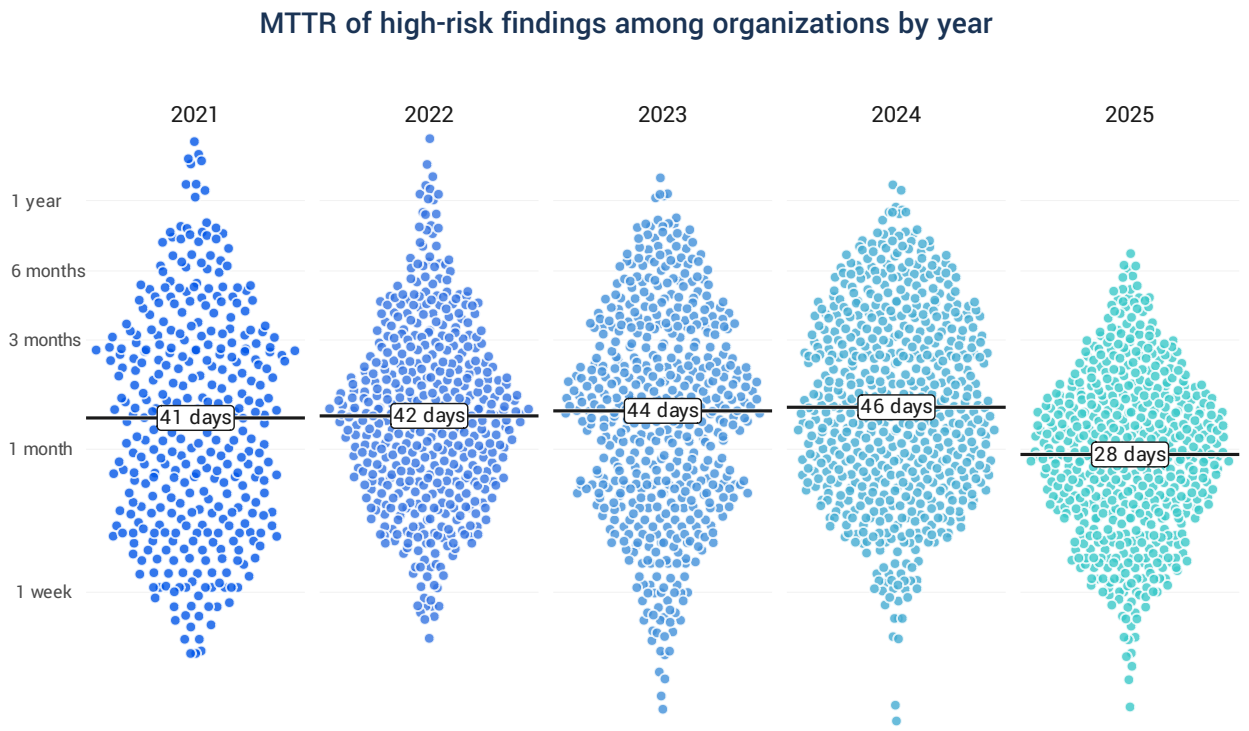

Figure 16

Source: Cobalt State of Pentesting Report 2026

Penetrating Insight

The median organization takes 39 days to resolve high-risk findings. That exceeds the SLA targets most organizations have set for themselves. Only the top-performing firms come close to meeting the most aggressive SLAs.

Half-life of findings

Like MTTR, the half-life of pentest findings also measures the speed of remediation. Unlike MTTR, half-life also accounts for unresolved findings to track the comprehensiveness of remediation. Half-life is derived from survival analysis, a technique often used to model the time until an event occurs (e.g., death, failure, resolution of findings).

Consider this: an organization that resolves 30% of its findings in five days will show an excellent MTTR, yet 70% of issues remain open to exploitation. Speed and completeness are both important. Organizations can't achieve a strong half-life by quickly fixing some findings while leaving many unresolved, or by resolving all findings at a glacial pace. That's why half-life tells the whole story of vulnerability remediation.

Figure 17 shows the overall survival timelines for pentest findings, and the effect of severity is clear. Critical findings have a half-life of just over a month. High and mid-severity findings persist a bit longer, with half-lives of about three and eight months, respectively. And low-severity findings? They never reach the halfway point within the timespan of this chart. They're understandably deferred in favor of more pressing security issues.

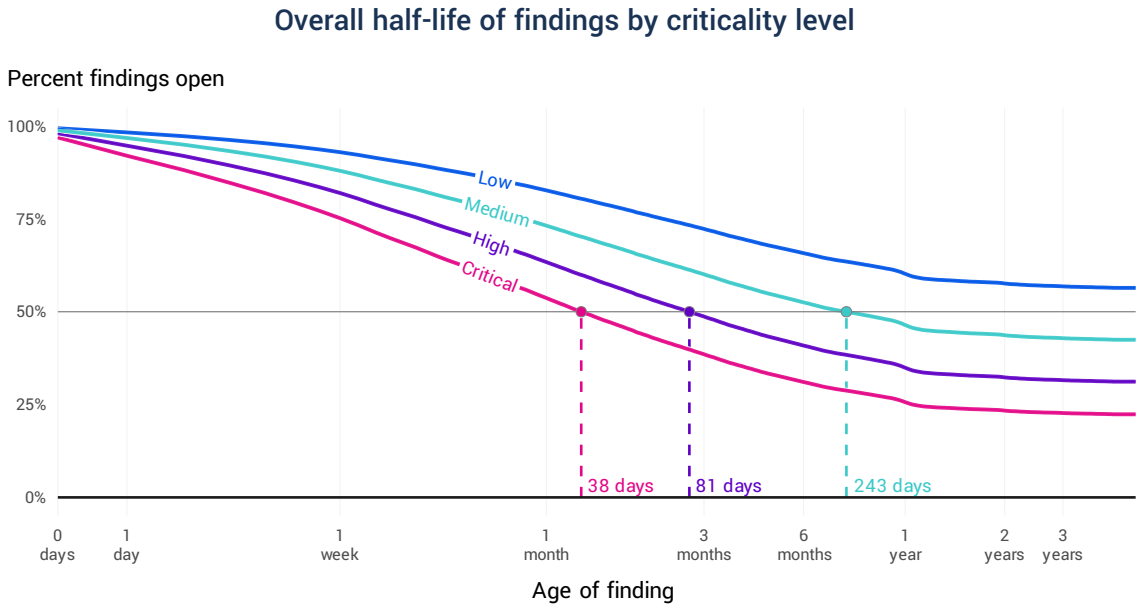

Figure 17

Source: Cobalt State of Pentesting Report 2026

Figure 17 applies survival analysis across all findings of all criticalities after accounting for organization-level differences. The more strategically useful view is produced by comparing organization-level half-life of high-risk findings to establish a performance range and identify what leading practice looks like (Figure 18).

We believe the half-life of high-risk findings is the single best measure of remediation performance. It rewards the critical trio of speed, completeness, and risk reduction. Nail all three, and your organization is far less likely to be one of the soft targets on attackers' hit list.

Let's establish some performance benchmarks in Figure 18. If you want to keep pace with the typical organization, aim for a half-life of 45 days—that's the median. Maybe that's good enough not to be easy prey. After all, you don't have to outrun the bear; you just have to outrun the slowest runner.

The performance gap between leading (top 10%) and lagging (bottom 10%) organizations in Figure 18 is striking. Top-performing organizations achieve a half-life of just 10 days for high-risk findings. Bottom-performing organizations have a half-life of 249 days—more than 8 months.

That 25x spread (10 days versus 249 days) isn't just a story about having plentiful vs. paltry resources. It reflects the presence or absence of a programmatic approach to penetration testing and remediation: conducting testing as part of a regular cadence, clear ownership of findings, integration with development and IT workflows, accountable SLAs for resolution rates, and a culture that treats pentest output as a live risk register rather than a periodic report to file away.

The pentest itself is table stakes. What separates a 10-day half-life from a 249-day one is everything that happens after the report lands. That distinction—between organizations that pentest and organizations that have built programs around pentesting—runs throughout this report. Half-life is perhaps the clearest single expression of it. Leaders close half their high-risk findings before most laggards have even triaged them.

Half-life of high-risk findings among organizations

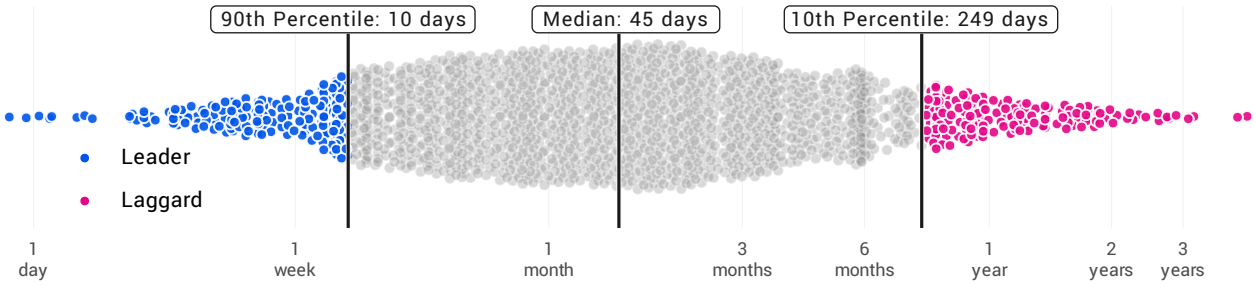

Figure 18

Source: Cobalt State of Pentesting Report 2026

Benchmarking sectors

From the metrics above, there's clearly a wide disparity among organizations when it comes to the prevalence and resolution of pentest findings. Does such performance variation exist among sectors? Let's find out as we close out this section with Figure 19. The columns represent the four benchmark metrics analyzed earlier.

- Percentage high-risk: Percent of pentest findings rated high-risk
- Percentage high-risk unresolved: Percent of high-risk findings unresolved
- High-risk MTTR: MTTR for high-risk findings
- High-risk half-life: Half-life of high-risk findings

Comparison of core performance metrics by sector

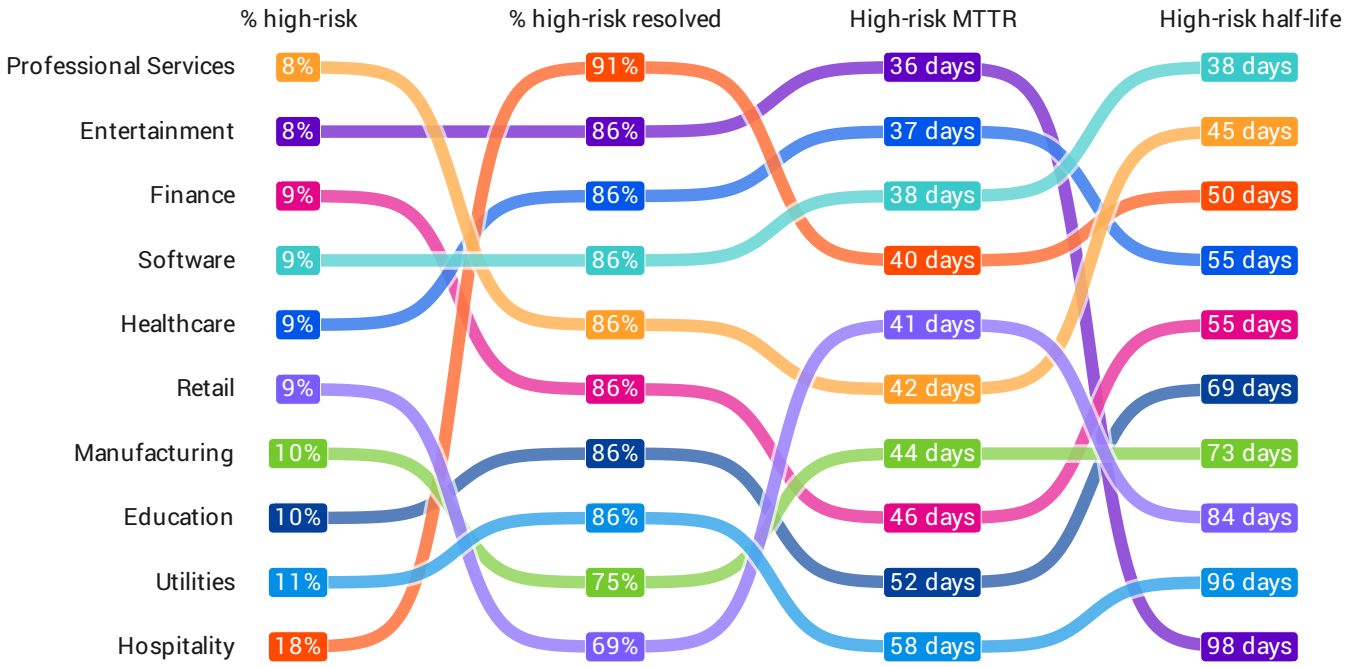

Figure 19

Source: Cobalt State of Pentesting Report 2026

Generally leading performers

- **Software:** Ranks in the top 4 for all resolution metrics. Perhaps an advantage of skilled development teams familiar with the code and able to address issues quickly.
- **Healthcare:** Stays in the top 5 across all resolution metrics, which is rather surprising, given the many challenges for IT and security teams in healthcare environments.

Generally lagging performers

- **Utilities:** Slowest MTTR and second-longest half-life for high-risk findings. Industrial infrastructure is built for resilience, not speed. Resolution requires caution and time.
- **Education:** Given the stereotype of this sector, the big surprise may be that it's not at the bottom. In fact, EDU almost creeps into the top half for two metrics.
- **Manufacturing:** Ranks among the bottom tier across the board. Similar to utilities, this likely ties to the industrial infrastructure that comprises the production floor.

Inconsistent performers

- **Hospitality:** If it weren't for the sector-leading rate of high-risk findings—which doesn't necessarily indicate poor performance—hospitality would be a surprise leader.
- **Retail:** Achieves above-average MTTR but drops to dead last for the comprehensiveness of resolution. Retail is the poster child for measuring half-life.
- **Finance:** Probably the most counter-intuitive of the bunch, because FinServ often sits atop for various security stats. Here, the sector floats in the middle of the pack.

Top Pentest Findings

Offensive security cuts through the noise. By replicating real-world attacker tactics, we identify proven vulnerabilities that directly compromise PII, financial data, and core business systems. With this lens, we asked our pentesters to weigh in on findings that provide the most leverage for attackers across three common asset types.

Web App and API

Cross-Site Scripting (XSS)	
Old Reliable	
Frequency	**9%**
High Risk	**32%**
Unresolved	**61%**

Missing Access Control	
The Side Door	
Frequency	**20%**
High Risk	**25%**
Unresolved	**58%**

1. Old Reliable: Cross-Site Scripting (XSS)

Despite being a 20-year-old problem, XSS is here for the long haul.

You may be asking "Doesn't [insert modern framework] solve this?" Sure, if you use it perfectly. But one `dangerouslySetInnerHTML` call later, and we're stealing session cookies. Pentesters love XSS because it's the ultimate bridge from a boring technical flaw to a high-value human target. While SQL injection goes after the database, XSS goes after the user—often seen as the weakest link.

It comes down to lateral movement at the app layer. In a recent assessment, we identified a Stored XSS vulnerability in a public feedback module. Our pentesters deployed a payload that persisted until an administrator viewed the logs; the script then exfiltrated the admin's session identifier to our listener, allowing us to hijack the session and gain full administrative access.

Beyond session hijacking, XSS enables pixel-perfect credential harvesting. By dropping a fake login overlay on a trusted URL, we can trick users into handing over passwords. With 61% of these findings going unresolved, XSS remains one of the most reliable and available entry points for attackers today.

2. The Side Door: Missing Access Control

Missing access control is pure efficiency for an attacker. It provides a toehold and sensitive data—like PHI or credit card numbers—without the need for complex privilege escalation or burning a novel Zero-Day. We look for the "broken doors" where a system fails to verify whether a user actually has the right to see the data they are requesting.

Insecure Direct Object Reference (IDOR) is the standout offender here. In a recent engagement at a major retail bank, we found an IDOR in their mortgage portal where the `account_id` was exposed in the URL. By simply incrementing that ID, our team bypassed all authorization checks to view sensitive financial statements and Social Security numbers of other applicants. It's a perfect example of how a simple oversight in object-level permissions can lead to a massive data breach.

Server-Side Injection	
The Master Key	
Frequency	**5%**
High Risk	**23%**
Unresolved	**60%**

Business Logic	
The Loophole	
Frequency	**2%**
High Risk	**9%**
Unresolved	**60%**

3. The Master Key: Server-Side Injection

Server-side injection hits like a sledgehammer. Whether it's SQLi or Command Injection, a single unsanitized input can flip a "harmless" search bar into a full-blown database shell. We don't just see these as bugs; they are fundamental failures of trust. If your app can't tell the difference between `user data` and `system commands`, you are in trouble.

Despite the risk, 60% of these findings languish in backlogs. Why? Because they often represent fundamental architectural failures rather than simple configuration toggles. In legacy codebases, you aren't just patching a line of code; you're rewriting how the application communicates to the database or the OS. One wrong move and you break the very business features that keep the lights on.

4. The Loophole: Business Logic

Business logic abuse is a playground of unintended consequences. These flaws exist in how a system is intended to function rather than how it is coded. Pentesters hunt for ways to manipulate the state of a transaction by violating the underlying assumptions of the developers—nuances that require a human to reason through "what should happen" versus "what is possible."

In a recent engagement for a major global distributor, we identified a critical credit processing flaw where submitting a negative balance triggered an automated refund. While the system returned a successful `200 OK` response, our pentesters demonstrated that an attacker could drain a corporate bank account by simply entering a minus sign to force a credit memo.

Cross-Site Request Forgery (CSRF)	
The Hijack	
Frequency	**2%**
High Risk	**11%**
Unresolved	**58%**

5. The Hijack: Cross-Site Request Forgery (CSRF)

CSRF is a high-impact stealth play that tricks a victim's browser into performing unauthorized, state-changing actions. While appearing in only 2% of findings, 58% go unresolved because remediation—such as enforcing `SameSite` attributes—often requires a major engineering lift in legacy code. In one engagement, we used a CSRF flaw to silently add an attacker-controlled SSH key to a developer's profile; to the server, it appeared as a standard update, but for the client, it resulted in a total compromise of their source code.

Mobile

Type	Frequency	High Risk	Unresolved
Authentication & Credential Vulnerabilities	0%	38%	75%
Missing Access Control	8%	20%	69%
Server-Side Injection	2%	18%	66%
Cross-Site Scripting (XSS)	3%	15%	67%
Business Logic	2%	11%	67%
Authorization & Privilege Escalation	0%	10%	90%
Authentication & Sessions	5%	5%	78%

Mobile vulnerabilities fall into two categories: local device issues and API-based flaws. While local weaknesses often require physical hardware access, API-based session flaws can be exploited remotely from anywhere in the world. Because mobile apps frequently share the same backend as web platforms, a hijacked mobile session is often a direct shortcut into your entire web infrastructure.

The fact that 78% of these vulnerabilities remain open represents a significant gap. By manually linking a mobile session token to a web administrative portal, our pentesters often chain separate medium findings into a single critical takeover. We consistently see Missing Access Control, XSS, and Server-Side Injection appearing where the mobile front-end connects to backend data.

For example, in a recent engagement for a global logistics firm, we identified a Session Token Reuse vulnerability. The mobile app generated a persistent token that lacked a “scope” limitation. Our team intercepted this token from mobile traffic and found it was also valid for the company’s administrative web portal. By injecting this token into a standard browser header, we bypassed Multi-Factor Authentication (MFA), transitioning from a low-level mobile user to a full web administrator with access to global shipping manifests and client PII.

Network

Type	Frequency	High Risk	Unresolved
Server Security Misconfiguration	33%	8%	62%
Active Directory (AD) & Domain Vulnerabilities	7%	36%	72%
Sensitive Data Exposure	7%	9%	69%
Network Misconfigurations & Exposures	3%	33%	63%
Authentication & Credential Vulnerabilities	3%	59%	45%
Privilege Escalation & Unauthorized Access	0.02%	100%	50%

While the industry fixates on the latest CVEs, our pentesters get most excited when they see the structural cracks that define modern internal networks: Active Directory (AD), Privilege Escalation, and Authentication flaws. These are the keys to the kingdom. Privilege Escalation and Unauthorized Access findings are always high-risk (100%), yet they remain unresolved half the time—a persistent vulnerability that bypasses the need for complex exploits by abusing how a network trusts its users.

The path to compromise often starts with legacy protocols like LLMNR or NetBIOS. When combined with disabled SMB signing—a classic misconfiguration—we can relay captured credentials to move laterally across the environment. Active Directory and Domain Vulnerabilities have surged in frequency by +7%, and with a 72% unresolved rate, they provide a near-guaranteed path to total control.

Ultimately, Authentication and Credential Vulnerabilities (59% are high-risk) allow for the elevation of privileges until a Domain Controller is reached. At that point, the entire identity layer is compromised. The fact that 72% of AD vulnerabilities and 45% of credential issues remain open represents a critical gap. To harden these environments, security teams must prioritize enforcing SMB signing and decommissioning legacy protocols.

CONCLUSION

The Offensive Security Leaders Quadrant

Ultimately, the State of Pentesting Report is a guide for organizations to benchmark their performance in closing security gaps. So for the first time, we opted to create an Offensive Security Leaders Quadrant that serves as a roadmap for security leaders and executives, by demonstrating that strategy and performance are intricately linked.

Based on our survey of 450 security executives and practitioners, we identified how organizational maturity—taking an ad hoc or programmatic approach to offensive security—directly impacts the ability to hit critical safety targets—establishing and hitting service level agreements (SLAs). With these two criteria, we identify four groups.

- Strategic Leaders: Organizations that pair a programmatic mindset with elite performance on resolving the highest-risk vulnerabilities within their SLAs. They set aggressive three-day SLAs and have the operational maturity to actually resolve 45% of critical findings within that window.
- Programmatic Ascendants: Teams that have moved toward continuous testing, but are still maturing the internal remediation workflows needed to handle findings and meet critical SLAs at scale.
- Compliant Accelerators: Teams that move fast to fix what is required to hit compliance goals, but are limited by a check-the-box mentality. These organizations often miss their SLA deadlines.
- Tactical Teams: Organizations trapped in an ad hoc testing cycle are laggards in performance against SLAs. These teams face the highest exposure risk.

By establishing this framework, we invite CISOs and security professionals to map where they are today, and consider what changes they might implement to improve outcomes.

The Offensive Security Leaders Quadrant

Source: Cobalt State of Pentesting Report 2026

Recommendations for Scaling a Strategic Offensive Security Program

Adopt a programmatic, continuous approach to pentesting

A programmatic strategy is a strong predictor of success in meeting SLAs and closing security gaps faster. Top performing organizations move away from ad-hoc testing or purely compliance-driven models in favor of a continuous calendar of testing, frequently leveraging pentesting as a service (PTaaS) to do so.

Proactively work to achieve executive and board alignment

Top-performing organizations align their objectives to business strategy, gaining board support and access to more resources. They ensure the board views offensive security as a necessary component of security posture and a strategic business enabler, rather than just another task. This strategic positioning makes it highly likely that the board will support resource and budget needs for offensive security programs.

Build strong cross-functional collaboration and accountability

A massive 25x performance gap exists between leaders (who have a 10-day half-life for remediating high-risk findings) and laggards (who take 249 days). Reduce your exposure window by establishing clear ownership of findings, integrating remediation into development and IT workflows, holding teams accountable to agreed-upon SLAs, and treating pentest findings as a live risk register.

Prioritize remediation based on business impact

Rather than prioritizing vulnerabilities that are easiest to fix, or to pass compliance audits, leading security teams focus their remediation efforts based on the potential business impact of the vulnerability and the risk of PII or sensitive data exposure.

Enforce strict third-party supply chain governance

While risks associated with third-party software are a leading IT concern, many organizations still deploy vendor software without requiring initial proof of security, creating a security debt that may end up being paid further down the chain by customers. Break the chain of vulnerability by requiring penetration test reports for all purchased software and tools before they are deployed.

Proactively test and monitor AI/LLM implementations

With AI-related security incidents—particularly those driven by unapproved shadow AI—becoming a leading threat, top performing organizations are moving toward programmatic testing of LLMs, rather than waiting to test reactively after an incident occurs. This involves pentesting AI, and partnering with pentesters who have specialized expertise in LLM.

What is the state of your pentesting?

The data in this year's State of Pentesting Report is clear: the 25x remediation gap—where top performers resolve high-risk findings in 10 days, while laggards take 249 days—is a strategic choice, not a resource constraint.

By transitioning from ad hoc, compliance-only testing to a continuous PTaaS model, organizations become 4.5x more likely to hit critical three-day SLAs, and close the dangerous disconnect between executive perception and practitioner reality.

In an era where AI-driven applications harbor high-risk vulnerabilities at 2.7x the rate of traditional software, maintaining a programmatic offensive security posture is the only way to transform security from a static report into a decisive business enabler.

Close your remediation gap.

Don't let high-risk findings linger for months. See how our human-led, AI-powered™ platform helps you move into the Strategic Leader quadrant.

TALK TO AN EXPERT

Research Methodology

This report analyzes two different datasets. The majority of analysis is based on data collected during Cobalt pentests. This is supplemented by insights collected via a survey by a third-party research firm, Emerald Research. These are further described below.

Pentesting methods

Cobalt pentesters follow specific methodologies depending on the type of test they're performing. By default, that includes industry-standard vulnerabilities from the Open Web Application Security Project (OWASP), which includes different top 10 lists for web, API, mobile, AI/LLM, and cloud systems.The Open Source Security Testing Methodology Manual (OSSTMM) is used for pentests of internal and external networks. More details on each of our pentesting methodologies can be found on the Cobalt website.

Pentesting data analysis

All penetration testing data analyzed in this report was collected by Cobalt pentesters. This spans over 16,500 pentests conducted on nearly 3,000 organizations over a 5-year period. Metadata from these pentests was exported from the Cobalt platform, sanitized to remove client-identifying and other sensitive details, and provided to Cyentia Institute for independent analysis. All statistics and charts featured in this report were produced and validated by Cyentia.

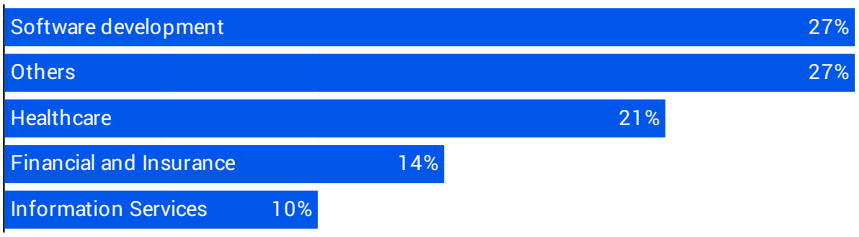

Source: Cobalt State of Pentesting Report 2026

Organizational metrics

We estimate organization-level metrics using statistical models that improve stability by accounting for repeated observations within organizations and differences between them (hierarchical or partial-pooling models).

Empirical Bayes beta-binomial models estimated organization-level rates (% high-risk, % high-risk resolved), reducing the influence of small sample sizes by shrinking noisy estimates toward the overall mean.

Mixed-effects count models and frailty-based survival models (random-effects Cox models) were used to estimate typical time-to-event (MTTR and half-life, respectively), the latter of which includes observations with censored outcomes in order to incorporate unresolved observations appropriately.

Survey sample

Cobalt contracted Emerald Research to administer a double-blind survey. About 1,682 people responded to the invitation, but over 73% were disqualified based on screener questions and various quality checks during and after the survey.

In the end, a total of 450 validated responses were collected. The sample consists of full-time information security leaders (50%) and practitioners (50%). These participants represent a range of industries and organization sizes.

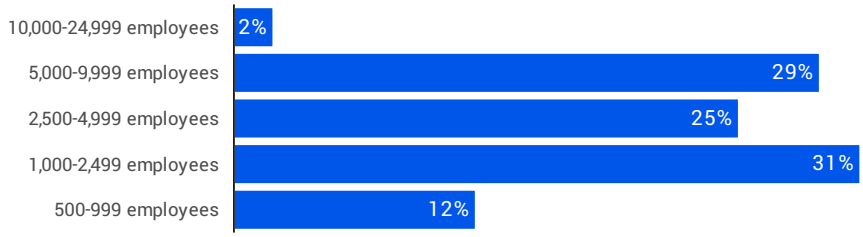

Source: Cobalt State of Pentesting Report 2026

About Cobalt and Cyentia

Cobalt is the pioneer in penetration testing as a service (PTaaS) and a leader in human-led, AI-powered offensive security™ services. We are focused on combining talent and technology with speed, scalability, and expertise. Thousands of customers and hundreds of partners rely on the Cobalt Offensive Security Platform, along with 500+ trusted security experts, to find and fix vulnerabilities across their environments. By enabling faster pentest launches, real-time collaboration with pentesters, and seamless integration with remediation workflows, we help organizations identify critical issues and accelerate risk mitigation so they can operate fearlessly and innovate securely.

The Cyentia Institute is a widely respected research and data science firm that works to advance cybersecurity knowledge and practice. It accomplishes that goal by collaborating with security companies to publish data-driven reports like this one and by providing analytic services that help organizations manage cyber risk.