



The state of AI pentesting in 2026

Adoption, trust, and the future of the pentest

intruder.io

October 2026

Introduction

AI is speeding up how software is built, which means more code, more changes, and more chances for vulnerabilities to reach production. It's speeding up how attackers find and exploit them too. Security testing has to keep pace with both.

For many organizations, that's a stretch. Most teams now ship code every day, but pentesting has long been something booked, scoped, and scheduled weeks in advance. Intruder's state of AI pentesting report, based on a survey of 500 security and engineering leaders across the US and UK, looks into how wide that gap has become.

Only 4% of organizations believe their current testing frequency gives them sufficient coverage, nearly half go more than a month between pentests, and more than two in five say they accept the risk and ship code without additional security checks in between. Automated tools such as DAST fill some of that gap, but many of the issues they flag still need to be validated manually.

As AI-assisted development puts more code into production, there are more of those findings to deal with, adding to the pressure on already stretched security teams. More than half spend six hours or more a week validating them, and fixes often take days or weeks to be confirmed.

That's starting to change. AI pentesting makes in-depth testing on demand possible, and adoption has moved quickly, with 92% of organizations having already tried it and 55% using it across their environment. Doubts about its quality aren't the main obstacle, with only a quarter seeing it as a top barrier. An approach that has been the norm for decades is giving way, with nine in ten agreeing that the traditional annual pentest is dead and continuous AI pentesting is now required to keep pace.

This report explores how that shift is playing out: how far testing has fallen behind development, what that costs teams, how widely AI pentesting has been adopted, how far teams trust it, and where this leaves manual testing.

Methodology

This report is based on a survey conducted by Censuswide in September 2026. Questions covered five core themes:

- How often organizations deploy code and pentest their applications
- How organizations cover the gap between pentests
- AI pentesting adoption, barriers, and trust
- How teams validate findings and confirm fixes
- The future role of manual pentesting

Job roles

- VP/Director of Cybersecurity (32%)
- CTO (23%)
- DevSecOps Lead (22%)
- CISO (11%)
- VP/Head of Engineering (8%)
- Chief Risk Officer (3%)

Company size

- 250–499 employees (12%)
- 500–999 (32%)
- 1,000–4,999 (44%)
- 5,000–10,000 (13%)

Geographies

- US (60%)
- UK (40%)

Sectors

- Technology
- Financial Services
- Manufacturing
- Professional Services
- E-commerce
- Healthcare
- SaaS
- Fintech

Security team size

- 2-5 (37%)
- 6-10 (52%)
- 11 or more (7%)
- 1 or fewer (3%)

Daily releases, periodic tests



68% of organizations deploy code daily, but nearly half go more than a month between pentests


Shipping daily is the norm

Most organizations now deploy code to production every day, with 68% deploying at least daily and 37% multiple times a day. The pace is higher in the US, where 74% deploy daily compared with 60% in the UK.


Testing still runs on a calendar

Running a pentest as often as code ships has never been realistic, largely because of the cost, so most organizations test periodically. Nearly half (47%) go more than a month between full pentests on their web applications and APIs, and just 1% test on every major release. That window between tests has always existed, but it's riskier now that AI is helping attackers find and exploit vulnerabilities faster. With most organizations shipping daily, a vulnerability introduced in one release can go weeks or months without a pentest. And these figures are for each organization's most frequently tested application, so the rest of their estate is tested less often still.

How often organizations deploy code to production



How often organizations run a full pentest on their most frequently tested application





42% of organizations accept risk and ship code without additional checks between pentests

Known risk, shipped anyway


Concern about the gap between pentests is widespread. 88% of security and engineering leaders running scheduled pentests are worried about unmapped risk introduced between tests, and they split almost evenly on how serious it is. Half call it one of their largest unmanaged cyber risks, with US leaders more likely to see it this way (51% vs 33% in the UK), while the other half describe it as a known gap they're covering with basic perimeter controls. Only 4% believe their current testing frequency gives them sufficient coverage.

That concern hasn't stopped many organizations from shipping code anyway. 42% say they accept the risk and ship code without additional security checks between pentests. The other ways teams cover the gap are spread evenly across automated scanning in the build pipeline, more frequent manual pentests, internal AppSec teams, and AI-driven pentesting, each used by around half. US organizations lean toward automated scanning (56% vs 45% in the UK), while UK organizations lean toward more frequent manual pentests (48% vs 54%).

The gap is widest for the midmarket

Half (50%) of organizations with 1,000 to 4,999 employees accept the risk and ship code without additional security checks between pentests, more than both smaller and large companies. They deploy just as often as the largest organizations, typically with smaller security teams, and are the least likely to pentest at least monthly (47%). It suggests this is where security teams are under the most pressure to keep up. That fits with Intruder's [Security Middle Child](#) report, where 42% of midmarket security teams described themselves as stretched, overwhelmed, or consistently behind.

Accept the risk and ship code without additional security checks between pentests, by company size



Tried by most, half use it at scale




92% of organizations have already tried or adopted AI pentesting

Adoption is already broad

55% of organizations say they use AI pentesting across their environment, and a further 36% have tried it but not yet adopted it at scale. Fewer than 1% have no plans to try it at all. For a category that has emerged in the past 24 months, that is fast uptake and a clear sign that security leaders and practitioners see AI testing as the way forward.

Has your organization adopted AI pentesting?



The largest organizations are furthest along

Organizations with 5,000 to 10,000 employees have moved furthest from trial to full use, with 71% using AI pentesting across their environment. Midmarket organizations have tried it at a similar rate but scaled it less so far, with 54% using it across their environment and 39% still trialing.

The US leads on adoption

US organizations are further ahead, with 64% using AI pentesting across their environment compared with 42% in the UK, where more organizations are still trialing it.

By sector, professional services (64%) and financial services (63%) lead, while manufacturing trails at 39%, and 18% of manufacturers haven't tried it at all, more than twice the overall rate. Financial services' high adoption may reflect the pressure it's under. It ships faster than most sectors, with 46% deploying multiple times a day (vs 37% overall), and handles some of the most sensitive data of any industry.

Confident in AI, cautious about access



Only a quarter see AI testing quality as a top concern


Privacy and safety top the list of barriers

Doubts about the quality of AI pentesting rank relatively low. Only 25% name its quality or depth compared with skilled human testers as one of their biggest barriers. The top barriers relate to how AI tools operate inside sensitive environments. 40% name data privacy, 36% safety on production systems, and 28% whether auditors or regulators will accept AI pentest reports. A further 24% point to a lack of internal expertise to evaluate or manage the tools.

More barriers in the midmarket


Midmarket organizations with 1,000 to 4,999 employees face the most obstacles. Only 4% say there are no major barriers, compared with 18% of organizations with 250 to 499 employees and 14% of the largest. They are also the most concerned about auditor acceptance (33%). Combined with typically leaner security teams to navigate and manage those risks, these barriers may help explain why they have tried AI pentesting at the same rate as the largest organizations but scaled it less.

The largest organizations are furthest along in adoption. They are the least likely to doubt AI testing quality (19% vs 26% in the midmarket), but they are the most concerned about privacy (46%). Wider use appears to ease doubts about what the tools can find, while concerns about what they can access remain.



Trust depends on transparency

When asked what would most damage their trust in AI pentesting, nearly half (46%) point to not being able to check the work: for 25%, that would be a lack of visibility into what code paths or endpoints were actually tested, and for 21%, a lack of reproducible, evidence-backed proof of exploitability. Hallucinated or false positive findings that waste developer time come close behind (23%), while missed vulnerabilities rank lowest (18%).



DevSecOps leads are the most skeptical

Concerns vary by role. CISOs are the most concerned about not being able to see what was actually tested, with 35% naming it as the biggest risk to trust, compared with 25% overall. DevSecOps leads, who are closest to the work and deal most directly with the findings security tools produce, are the most skeptical. They are more concerned than average about false positives (29% vs 23%) and missed vulnerabilities (26% vs 18%), and only 5% say nothing would damage their trust, compared with 19% of CTOs.

What would most damage trust	CISOs	CTOs	DevSecOps leads
Lack of visibility into what was tested	35%	22%	25%
False-positive findings	24%	17%	29%
Missed vulnerabilities	11%	15%	26%
Nothing would	13%	19%	5%

Most of these concerns can be checked before a tool is given access to your code or systems. [We've put together the questions worth asking vendors.](#)

Hours on triage, weeks to verify



56% of teams spend six or more hours a week triaging false positives


This is where AI-enabled validation makes a difference. It checks findings from security tooling before they reach the team, confirming which are real and exploitable. That means fewer false positives to triage, less time lost to dismissing noise, and more of the week spent fixing issues that matter.

Most of a working day spent validating, every week

Security tooling has long had a reputation for creating more work for already stretched teams. Every finding it surfaces has to be checked, and without proof of exploitability, that work is done by hand.


It's proving to be a major burden. The biggest remediation bottleneck is working out whether a finding is real or a false positive (44%). 56% of teams spend six hours or more a week triaging, validating, or dismissing false positives and unexploitable findings, and 21% spend 11 hours or more, more than a full working day. US teams spend considerably more time on it, and are more than twice as likely as UK teams to spend 11 hours or more a week (27% vs 13%).

Hours a week spent validating findings



Fixed, but not confirmed

Once a fix ships, confirming it worked often takes days or weeks. Fewer than two in five organizations independently retest a fix in production as soon as it's merged or released, while nearly half take up to two weeks. Midmarket organizations are slower than the largest, with 35% retesting immediately compared with 56%. In the meantime, there's no way to know the vulnerability has actually been closed. AI pentesting is designed to shorten that wait too, by retesting a fix as soon as it ships.



A funeral for the annual pentest



90% agree the traditional annual pentest is dead


Few will miss the annual pentest

90% of security and engineering leaders agree that the traditional annual pentest is dead and that keeping up with how often teams now ship requires continuous, AI-driven testing. Half (50%) strongly agree, and only 1% disagree. This reflects a shift that has already happened, with only 6% of organizations relying on a single annual pentest. US leaders feel it more strongly, with 60% strongly agreeing compared with 35% in the UK.

CTOs and security leaders are the most convinced

Views vary by role. Those who typically hold the budget for pentesting are the most likely to strongly agree that the traditional annual pentest is dead, including 65% of CTOs, 52% of VPs and Directors of Cybersecurity, and 45% of CISOs. DevSecOps leads (40%) and VPs of Engineering (27%) are less likely to strongly agree, which may reflect a different vantage point, closer to running tests and acting on findings than to procuring them.

Strongly agree the traditional annual pentest is dead, by role




A smaller role for manual testing

Most (54%) expect AI to handle the bulk of testing within three to five years, either largely replacing manual pentesting (28%) or leaving it for specialized work such as red teaming and social engineering (25%). Only one in six (16%) expect manual pentests to remain a primary method of security validation. Among the largest organizations, which are furthest along with AI pentesting, a third (33%) see a specialist role for humans, and only around one in ten (11%) expect manual testing to stay primary.

Manual pentests kept for auditors

28% expect manual pentests to be kept to satisfy compliance or audit requirements, regardless of their technical value, which means paying for a test mainly because an auditor asks for it. Some frameworks, such as SOC 2 and ISO 27001, focus on evidence of testing rather than how it's done, and AI pentest reports are already used to meet them. Guidance from auditors and regulators is still inconsistent and evolving, so while it settles, many organizations are likely to keep a manual pentest in the mix for compliance.

The role of manual pentests in three to five years



A priority without a definition




96% say continuous pentesting is a priority, but no single definition is chosen by more than 35%

Continuous pentesting is almost everyone's priority

Almost every organization calls continuous pentesting a priority (96%), and more than a third (36%) say they've already achieved it, with US organizations far more likely to say so (47% vs 19% in the UK). A further 44% are actively investing in it, while 16% say it's a priority but limited investment or execution is holding them back.

No agreed definition

The industry has yet to settle on what continuous pentesting means. Asked which of four definitions best matched their understanding, respondents split widely, with no single answer chosen by more than 35%. The most common is a combination of always-on attack surface monitoring and vulnerability management with regular manual validation (35%). With the term applied to very different approaches across the market, "achieved" is likely to mean different things to different organizations.



Testing at the speed of shipping

Until now, teams have relied on two kinds of testing, each with a trade-off. Traditional vulnerability scanning and DAST run often but are generic by design, so they miss flaws specific to how an application works. A pentest goes deep, but it's a snapshot of how the application looked weeks ago. AI pentesting brings the depth of a pentest at the frequency of a scan.

That changes what's possible across the whole testing cycle. Releases can be tested as they ship, instead of waiting for the next scheduled pentest. With findings validated before they're reported, teams get time back and can focus their effort on issues that are real, and fixes can be retested straight away.

Keeping pace now depends on testing like this, and a few steps make it easier to adopt with confidence. Start in a staging environment that mirrors production. Run an AI pentest alongside your next manual pentest and compare the results. Look for findings that come with proof-of-concept code and a record of what was tested. And check how your code and data will be handled before you give it access.


With AI speeding up how software is built and how it's attacked, AI pentesting makes it possible to keep pace with both.



About Intruder

Intruder's continuous exposure management platform helps security, IT, and engineering teams stop breaches before they start. By unifying AI pentesting, attack surface monitoring, cloud security, and vulnerability management in one intuitive platform, Intruder gives stretched teams an always-on security source of truth. Our approach focuses on continuous automated scanning using deep security expertise and agentic solutions to ensure that the findings we deliver are accurate, prioritized by real-world risk, and ready to act on.

Founded in 2015 by Chris Wallis, a former ethical hacker turned corporate blue teamer, Intruder is now protecting over 3,000 companies worldwide. Intruder has been awarded multiple accolades, was selected for GCHQ's Cyber Accelerator, included on Deloitte's Tech Fast 50 2023 list as the fastest-growing cybersecurity company in the UK and was named in G2's 2026 Best Software Awards.



Start a free trial or book a call with one of our experts at intruder.io



Read our reviews on [G2.com](https://g2.com)