Stingrai

AUGUST 2026

The State of Penetration Testing 2026

An analysis of 1,206 verified findings from 55 penetration tests

1,206	55	67.2%	0.74%
verified findings	penetration tests	High or Critical	false positives

THE YEAR IN NUMBERS

Data from Stingrai's penetration testing platform, extracted on 29 August 2026.

1,216

findings from 55 penetration tests across different industries

67.2%

of verified findings were High or Critical: 810 of 1,206

93%

of penetration tests surfaced at least one High or Critical finding; 69% surfaced a Critical

8

median findings per test; the middle half of engagements landed between 5 and 15

0.74%

false-positive rate: 9 findings declined at two-stage review

26 days

median time to resolve where tracked; the slowest tenth ran past 121

92% vs 54%

share of findings rated High or Critical in internal network testing against web application testing. What a test finds depends on what you test, and this report is segmented accordingly.

KEY TAKEAWAYS

The single most important number in this report is that there is no single number.

What a test finds depends on what you test.

Outdated software is 56.9% of internal network findings and a rounding error in web application work, where authentication and session lead at 28.1%. A pooled benchmark describes an engagement nobody bought.

A pentest is two kinds of work in roughly equal measure.

37.2% of findings fall into classes a scanner could in principle flag. 35.2% fall into classes no signature can represent: authentication, access control, IDOR, business logic.

The severe-findings headline is largely an internal network effect.

Outdated software is the largest block of High and Critical findings at 28.6%, driven by internal network estates rather than application code.

One severe finding in three involves logic or identity no scanner models.

Authentication, authorization and business logic together are 34.4% of High and Critical findings.

Critical findings get fixed fast. High severity is where fixes stall.

Median 10.5 days to resolve a Critical finding against 38.0 for High, the opposite of the usual assumption.

What earns a Critical is not what earns a High.

High findings are led by outdated software at 32.8%. Critical findings are led by authentication at 28.5%, and an injection finding was rated Critical 56.5% of the time, the highest of any class.

TWO THIRDS OF FINDINGS ARE HIGH OR CRITICAL

Only 5 findings in the corpus were logged as Informational: the bottom of the distribution is truncated by reporting practice.

67.2%

of verified findings were High or Critical: 810 of 1,206.

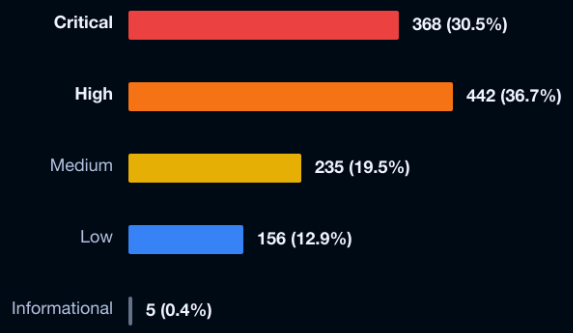

Source: Stingrai penetration testing platform data, October 2024 to August 2026.

WHAT EARNS A CRITICAL: IDENTITY AND INJECTION, NOT PATCH DEBT

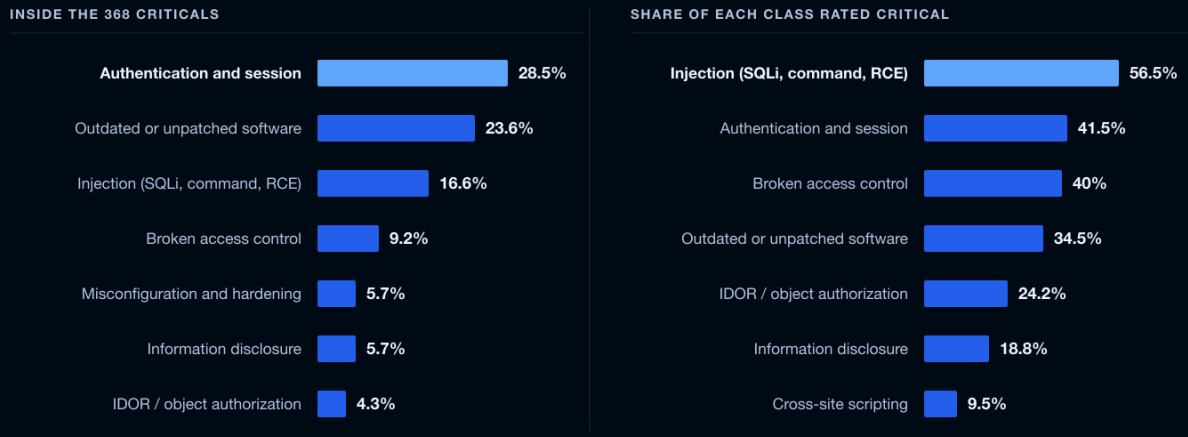

Source: Stingrai penetration testing platform data, October 2024 to August 2026.

WHAT PENETRATION TESTS ACTUALLY FIND

Authentication and session leads outdated software by exactly one finding.

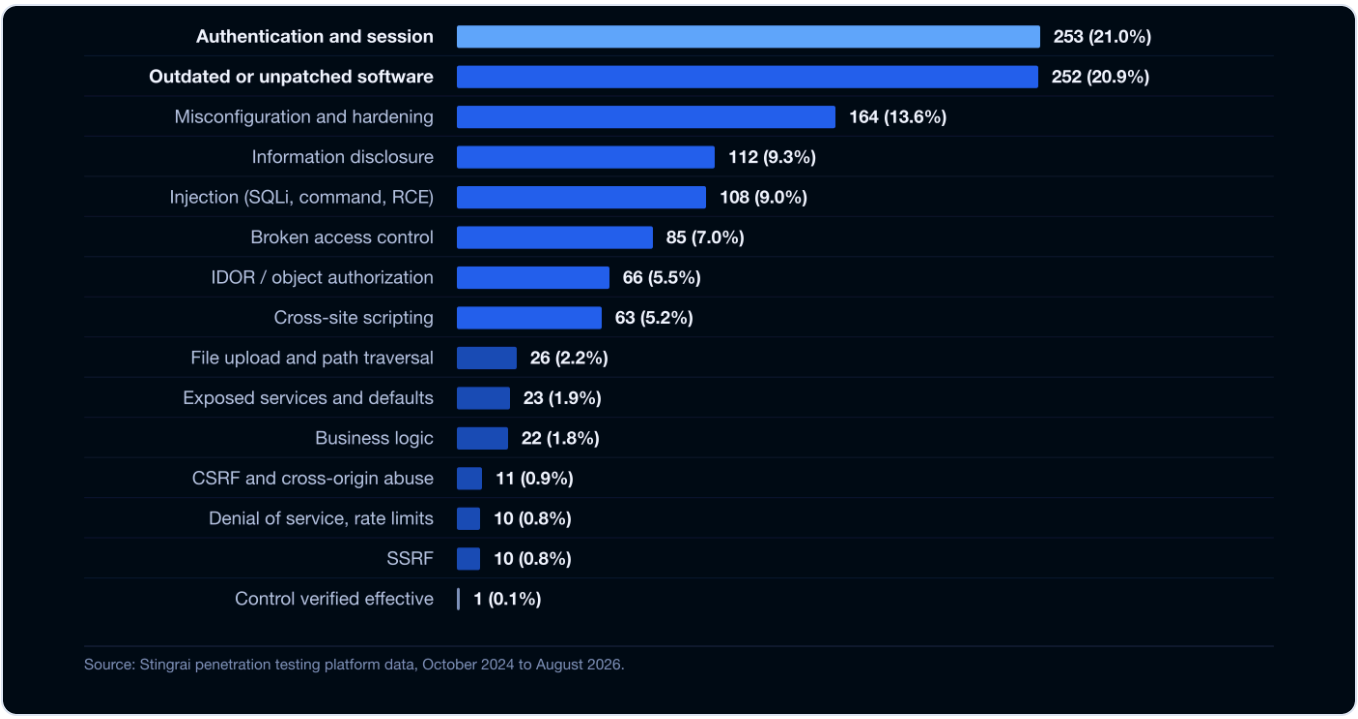

TWO KINDS OF WORK, IN ALMOST EQUAL MEASURE

37.2%

scanner-shaped classes: outdated components, hardening gaps, exposed services. They have signatures and can be checked on every deploy.

35.2%

context-dependent classes no signature can represent: authentication, access control, object-level authorization, business logic.

THERE IS NO AVERAGE PENTEST

Internal network engagements run 92% High or Critical; web application engagements run 54%.

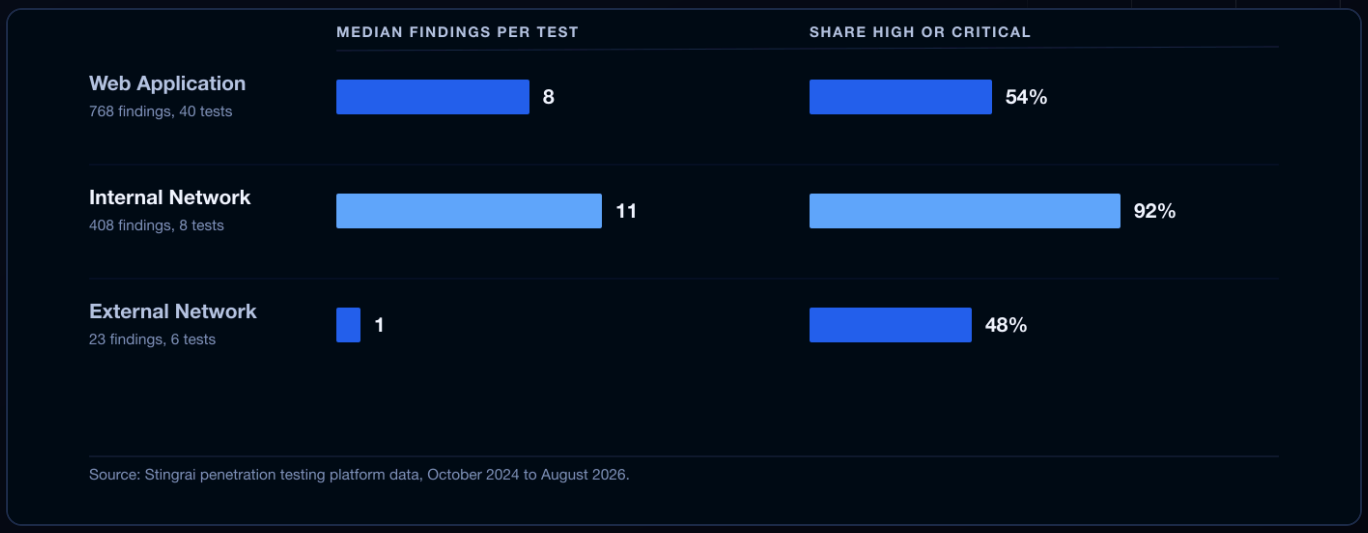

Source: Stingrai penetration testing platform data, October 2024 to August 2026.

THE SAME CORPUS, TWO COMPLETELY DIFFERENT PICTURES

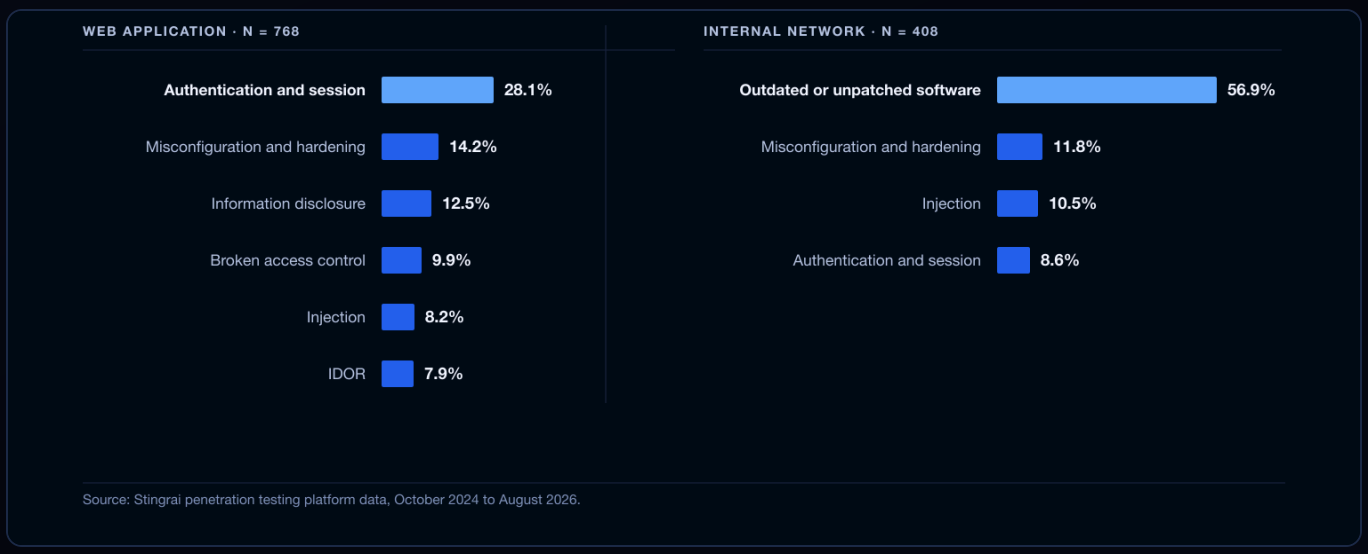

Source: Stingrai penetration testing platform data, October 2024 to August 2026.

WHAT ONE TEST PRODUCES

Median 8 verified findings per test. One engagement produced 223 and drags the mean to 21.9.

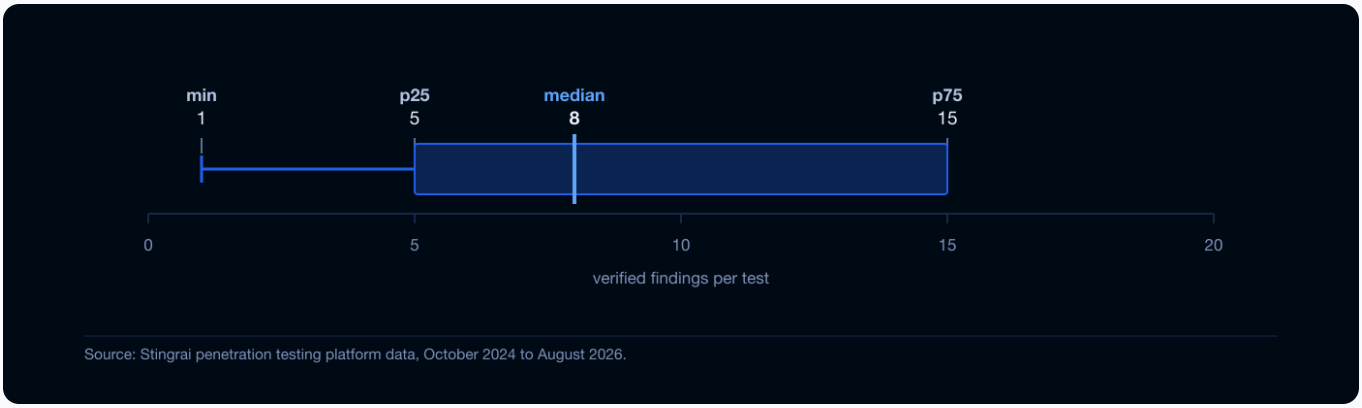

EVERY FINDING IS PROVEN BEFORE IT IS REPORTED

0.74%

false-positive rate: 9 of 1,216 logged findings declined at review, measured over everything logged, not a curated subset.

Each finding carries reproduction evidence and passes two stages of review, a team lead and then the engagement partner, before a client sees it. That discipline is what separates a verified finding from scanner output.

HOW LONG A FIX TAKES

Critical findings close at a median of 10.5 days. High findings queue with everything else and take 38.0.

26 days
median time to resolve

121 days
90th percentile

Source: Stingrai penetration testing platform data, October 2024 to August 2026.

Watch the High queue, not just the Critical queue.

An SLA that treats High as "urgent but not an incident" is where exposure accumulates, and a 30-day retest window never sees the tail.

WHAT THIS MEANS FOR DEFENDERS

01

Benchmark against your test type, not a pooled average.

Internal network: median 11 findings at 92% severe. Web application: 8 at 54%. Planning either off a blended figure gets both wrong.

02

Split the budget the way the findings split.

For internal networks, patch and lifecycle management is the programme. For web applications, the findings concentrate in authentication, access control and IDOR.

03

Run the scanner-shaped classes continuously, not annually.

They have signatures. Discovering them at pentest cadence is paying human rates for work that does not need human judgement.

04

Watch the High queue, not just the Critical queue.

10.5 days median for Critical findings, 38.0 for High. The tier below the incident line is where exposure accumulates.

05

Size the retest window against the p90, not the median.

Half of fixes land inside 26 days; the slowest tenth runs past 121. Retesting inside the engagement makes the tail visible.

METHODOLOGY, IN BRIEF

Source.

Stingrai's penetration testing platform, the system of record for engagements, findings and remediation, extracted on 29 August 2026.

Corpus.

1,216 findings logged across 55 penetration tests, after excluding internal demonstration engagements and test records. 9 declined as false positives. All distributions computed over the 1,206 verified findings.

Classification.

Every finding was classified by rule-based analysis of its title, applied uniformly across the corpus. The method reads titles only, which makes it conservative.

Limits.

One firm's book of work, not an industry survey. Resolution times cover the 163 findings with tracked timing. Mobile contributed one test and is counted but not segmented. The full methodology and limitations are open online.

About Stingrai.

Stingrai is a CREST-accredited penetration testing provider. Engagements combine Snipe, our autonomous web application pentesting agent, with human penetration testers, and retesting is included in every engagement.

Cite this report as: Stingrai, The State of Penetration Testing 2026. Full report, methodology and limitations: stingrai.io/blog/state-of-penetration-testing-2026. Data questions: hello@stingrai.io

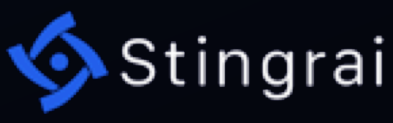
Stingrai