

State of Pentesting

in Financial Services and Insurance Industries

2026



The data tells the story


The financial services and insurance industries are disciplined when it comes to pentesting, with one blind spot. While the financial services and insurance industries test their AI applications more than any other sector and have the fewest AI incidents, they are planning to expand red teaming the least, and their remediation time is only mid-pack.

Examined together, these two industries run security like the regulated sectors they are: programmatic, measured, and confident. This state of pentesting snapshot combines five years of Cobalt pentest findings with our 2026 survey of security leaders and practitioners, to profile how these sectors find, fix, and think about security risk. Their teams test AI more than anyone and report the lowest AI incident rate of any industry. Yet underneath that confidence, remediation is only average on the truer half-life measure. And despite testing the most, these two sectors are the least likely to add the one practice built to find what routine testing misses: red teaming.


Sources: Cobalt State of Pentesting Report 2026 • 2026 survey of 455 security leaders and practitioners (Emerald Research).
Data from financial services and insurance industries are combined due to smaller sample sizes (n=65).

The high-risk remediation funnel


Every high-risk pentest finding runs this gauntlet, from discovery to closure. The sector resolves a high share of findings, but its half-life sits in the middle of the pack.



test their AI applications programmatically
(highest of any sector)



have had an AI or LLM security incident
(lowest of any sector)




plan to increase red team operations
(lowest of any sector)


Where the financial and insurance industry stands: remediation half-life by sector

The sector sits mid-pack on half-life (55 days), tied with the healthcare industry and well behind the leaders. Its mean time to remediate (MTTR) is also middling at 46 days (seventh of 10), so unlike some fast-fixing sectors, finance and insurance is not quick on either measure. This analysis leads with half-life because it is the fuller view: it captures the time to close half of all high-risk findings, including those still open, whereas MTTR counts only the issues already fixed. For a sector this disciplined about testing, an average remediation record is the first sign that discovery is outpacing the work of closing findings out.

The time to remediate half of all high-risk findings (half-life). Lower is better.



Here is the same half-life ranking on a single scale. Each sector is a colored dot matching the bar above. Note that finance (financial services and insurance industries) sits in the top half of industries, but trails the leader (software) by 17 days.



Confidence, mostly earned

The financial services and insurance industries rate their own security posture highly, and unlike some sectors, the testing discipline behind that confidence is real. Security teams test AI more than any other sector and report the fewest incidents. The table below pairs survey findings with the pentest data that supports or complicates survey participants’ perceptions. The perception fails to meet reality in remediation and red teaming, where the record of these industries is only average.


THE CLAIM (SURVEY)	THE REALITY (PENTEST DATA)
74% see AI as a tool, not a threat	9% have had an AI incident, the lowest of any sector
48% feel well-equipped to address AI security	48% test AI programmatically, the highest of any sector
62% run a programmatic pentesting cadence	55 days half-life, only mid-pack, and 46 days MTTR (7th)

AI in the financial and insurance industry


The survey data shows what these sectors worry about in AI applications, and how they are testing AI. The two industries’ AI concerns track their regulatory reality. Data leakage leads among concerns, as it does almost everywhere, but the standout is model bias: financial services and insurance industry security teams name it as a concern more than any other industry, reflecting the legal and reputational stakes of biased automated decisions in lending, underwriting, and claims.

Accountability leans toward the CISO (72%), though less sharply than in the healthcare industry. Engineering, product, and the data and AI teams carry meaningful shares too, so responsibility is more distributed than in some sectors.

Top AI security concerns named by financial and insurance security teams



Who gets blamed if AI causes an incident



ZOOM OUT: AI findings across industries

AI is the hardest asset class to secure across every industry. The financial services and insurance sectors test it well, but the AI resolution problem persists. [Read the AI and Pentesting Pulse Report 2026](#) for the full cross-industry picture and where autonomous pentesting is headed.

32%

of AI findings are high-risk
(vs 12% across all assets)

38%

of high-risk AI findings get
resolved (lowest of any
pentest type)

Recommendations for a programmatic approach to pentesting



Add red teaming to match your testing

These industries test AI more than anyone but plan to expand red teaming the least (17%). Programmatic testing confirms known controls, red teaming finds the unknown paths.



Keep the programmatic edge, extend it to new AI surfaces

A 48% programmatic AI testing rate is a genuine advantage. Make sure every new model, integration, and agent enters that same cadence rather than shipping ahead of it.



Manage to half-life, not just incident counts

A low incident rate is reassuring, but the mid-pack 55-day half-life shows findings still linger. Track how long it takes to clear half of all high-risk findings, and drive it down



Distribute AI accountability deliberately

Blame leans to the CISO (72%) but engineering and data teams share responsibility. Define who owns AI security decisions before an incident forces the question.



Treat model bias as a security and compliance issue

Model bias is a much higher concern (37%) in financial services and insurance. Test AI systems for biased or manipulable outputs in lending, underwriting, and claims, not just for classic vulnerabilities.

See where your AI and application security actually stands

Cobalt runs ~5,000 pentests a year across a broad range of industries. Get a demo to see how your remediation half-life and coverage compare.

[GET A DEMO](#)