

Cait: Cacilian AI Pentester

Continuous, AI-assisted pentesting

Cait is an autonomous AI pentester that thinks before it attacks and proves every finding. It explores how your web apps and APIs actually behave, plans targeted attacks, and reports back only exploit-validated, audit-grade issues with real HTTP evidence. Instead of relying on noisy scanners or waiting for the next manual engagement, teams get recurring, credible insight that reflects how their applications really work.

The problem

Traditional scanners are fast but shallow. They fire static payload lists, miss context and logic, and flood teams with noisy, low-value findings that still need manual verification.

Manual pen tests are deep but episodic. Every engagement needs scoping calls, scheduling, and expensive human time, so coverage is limited and reports go stale quickly as products change. The result is a gap between releases, audits, and real-world risk: critical apps change faster than they can be tested, while security teams struggle to produce evidence that stands up to auditors and customers.

How Cait works



Context-first exploration

Cait starts by exploring your application like a human pen tester:

- Maps and classifies flows, states, and roles instead of blindly crawling URLs.
- Builds an internal model of how the app responds and where sensitive actions live.
- Uses this understanding to choose attack paths instead of spraying generic payloads.



Targeted, adaptive attacks

Based on what it learns, Cait:

- Launches focused attacks for issues like SQLi, XSS, IDOR, SSRF, command injection, and common auth/authorization flaws.
- Adapts payloads and strategy in real time based on actual responses.
- Covers more flows than a human could test in the same timeframe, with a consistent methodology.



Exploit-validated findings only

Every “high-quality finding” must pass a validation pipeline:

- 1-3 reproducible HTTP request/response pairs with highlighted evidence.
- Clear impact description and remediation guidance.
- CVSS scoring so engineering and security can prioritize quickly.
- If Cait can’t prove exploitability, it doesn’t ship the finding.

Built to extend human testers

Cait is designed as the always-on layer of a pentest program:

- Runs with consistent checks and methodology.
- Leaves space for human testers to handle social engineering, mobile/binary testing, deep business-logic flaws, and bespoke assessments.
- Human testers can step in via optional add-ons, using Cait’s exploration history as a starting point when more complex work is required.

What you can expect

- Reports in as little as **2-3 hours**
- **Recurring, context-aware pentests** that reflect how your apps really behave.
- Exploit-validated findings with **concrete HTTP-level evidence**, not scanner noise.
- Faster triage and remediation with **clear impact, CVSS scores, and guidance**.
- **Consistent coverage** between manual engagements, without repeated scoping and project overhead.
- Evidence that can be reused for **SOC 2, ISO 27001, PCI, HIPAA**, and customer due-diligence requests.