HADRIAN

NOVA

AGENTIC PENTESTING THAT VALIDATES YOUR EXTERNAL EXPOSURES

THE AI THAT HACKS YOU SHOULD BE ON YOUR SIDE

Hadrian Nova is an agentic pentesting engine built on Hadrian's external exposure management platform. It runs structured offensive security tests against your external attack surface on demand, with validated findings delivered within hours. Nova doesn't detect theoretical risks. It exploits real vulnerabilities, chains attack paths, and delivers confirmed findings with reproduction steps and remediation guidance. The full penetration test report is validated by Hadrian's offensive security team, and structured for both technical teams and as compliance evidence.

5X	80%	100HRS
ROI COMPARED TO MANUAL PENTESTING	FASTER TIME TO RESOLUTION	EQUIVALENT HUMAN HOURS SAVED

CHALLENGE

Traditional penetration testing hasn't kept pace with the attack surfaces it's supposed to protect. Weeks of scoping and scheduling, high per-engagement costs, and fixed scope limit most organizations to intermittent testing, while their external exposure evolves daily. By the time results arrive, they are often outdated due to environmental changes, inconsistent across testers, and disconnected from the remediation workflows that need to act on them.

SOLUTION

Nova runs a structured hierarchy of specialized AI agents, each expert in a distinct attack class, coordinated by Hadrian's AI Orchestrator that reads the application or API's structure and assigns work to the right specialist agents. After the initial testing, the Orchestrator reviews every finding and unresolved lead, dispatching targeted follow-up tasks where there are signs of weakness. The Orchestrator doesn't follow a fixed playbook. It monitors agent progress in real time, concentrating effort on the most promising attack vectors and deprioritizing dead ends, so every test delivers the highest-impact findings first.

Nova's agents don't work alone. Hadrian's offensive security specialists steer the investigation, provide input that sharpens agent targeting, and validate every finding. The result: zero false positives and a compliance-ready penetration test report, executive summary for leadership, per-finding reproduction steps for engineering, and full methodology documentation for auditors.

BENEFITS

- **ZERO FALSE POSITIVES**
  Every finding is validated by Hadrian's offensive security analysts before delivery.
- **ON-DEMAND SCHEDULING.**
  Run before a release, ahead of an audit, after a configuration change. No lead time, no vendor coordination.
- **80% FASTER TIME TO RESOLUTION.**
  Findings include reproduction steps, exploitation evidence, and remediation guidance — ready to act on immediately.
- **TOTAL SCOPE CONTROL.**
  Define the target, exclude sensitive areas, provide application context and test focus areas, all without renegotiating an SOW.
- **10X VISIBILITY OF CRITICAL VULNERABILITIES.**
  Agents test the full defined scope without time pressure, surfacing exposures that time-boxed manual tests miss.
- **ADAPTIVE COVERAGE.**
  Nova evolves with attacker TTPs. Coverage deepens over time, not just repeats.
- **COMPLIANCE-READY.**
  Report satisfies SOC 2 (CC7.1), ISO 27001 (A.8.8), NIS2 (Article 21), and DORA (Article 25) audit requirements.

INTRODUCING NOVA: AGENTIC PENTESTING

Nova is an agentic pentesting engine that runs offensive security tests against web applications and APIs on your external attack surface. Nova covers the OWASP Top 10 from an external testing perspective, including authentication bypass, injection attacks, access control, SSRF, XSS, CSRF, and server-side exploitation. Multi-account testing validates cross-user access and privilege escalation automatically. Categories that require source code access or internal infrastructure fall outside the scope of any external penetration test.

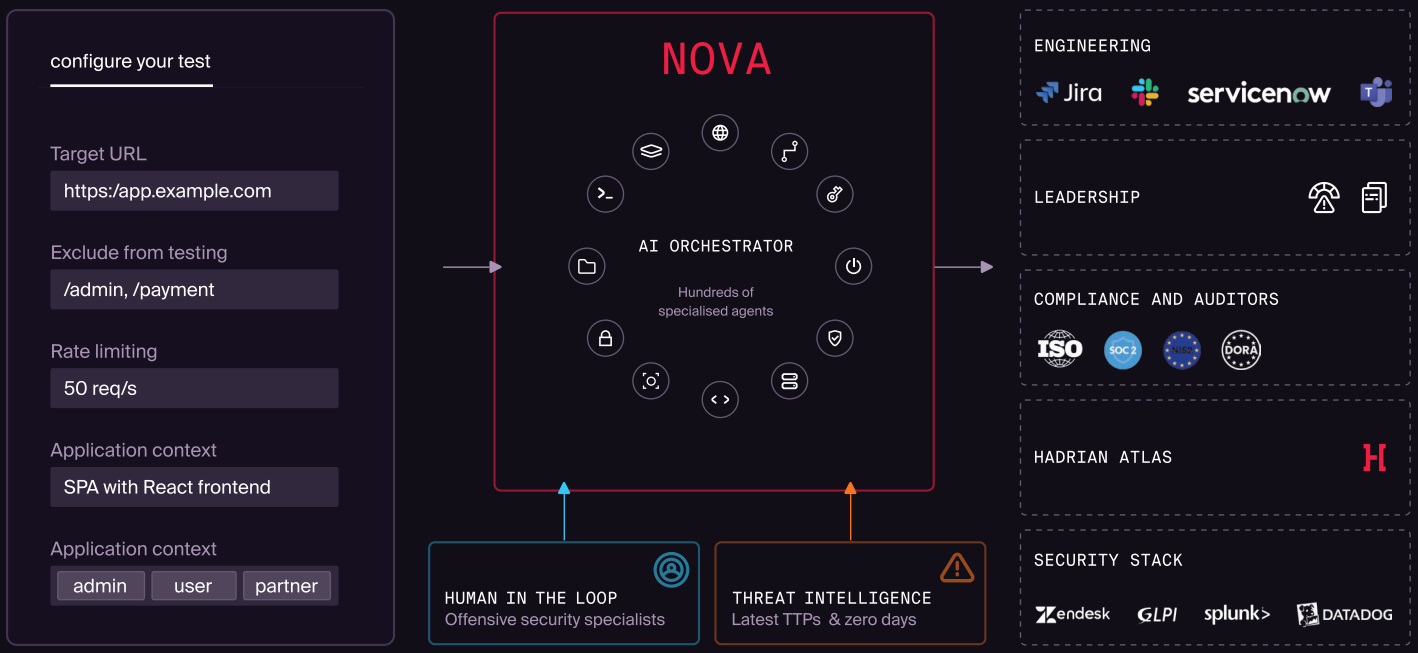

KEY CAPABILITIES

- **MULTI-ACCOUNT TESTING.**

Provide credentials for multiple roles. Agents test cross-user access and privilege escalation automatically.

- **SCOPE EXCLUSIONS.**

Define areas that should not be tested via free text — no rigid scoping templates.

- **CONTEXT-AWARE RECONNAISSANCE.**

Discovery is seeded with Hadrian's existing asset intelligence. Endpoints are grouped by attack surface for targeted testing.

- **COMPLIANCE-READY REPORTING.**

Executive summary, methodology, risk ratings, reproduction steps, and remediation guidance. Satisfies SOC 2, ISO 27001, NIS2, and DORA.

- **HIERARCHICAL AGENT ARCHITECTURE.**

Specialized agents per attack class, orchestrated for coverage and reliability.

- **HACKER-IN-THE-LOOP.**

Elite offensive security professionals steer, validate, and sign off.

RECOGNISED BY LEADING ANALYSTS

Gartner Peer Insights 4.9/5 ★

GIGAOM 2025

FROST & SULLIVAN BEST PRACTICES AWARDS

COMPLIANT WITH

AICPA SOC

ISO 27001

ABOUT HADRIAN

Hadrian is an external exposure management provider that pioneered the AI attacker's perspective approach. Its agentic engine offers frictionless always-on discovery, validation, and mobilization of an organization's most critical cyber risks. Trained by elite hackers with top offensive knowledge, Hadrian adapts the organization's unique environment to continuously probe, discover and validate the risks that attackers can actually exploit.

TO LEARN MORE: HADRIAN.IO

BOOK A DEMO