EC-Council Investee Company FireCompass' AI Pentest Agent Reaches HackerOne's Top 3 on a $5,000-a-Month Budget

Live results show agentic AI can uncover real vulnerabilities at scale and accelerate the shift toward continuous security validation

For EC-Council, the milestone provides tangible validation of an ecosystem strategy focused on combining advanced AI with skilled cybersecurity professionals, disciplined engineering, and accountable human oversight.

TAMPA, FL, UNITED STATES, August 4, 2026 /EINPresswire.com/ -- Autonomous penetration testing has moved from controlled experimentation to measurable performance in live security environments. FireCompass, an EC-Council investee company, has taken its AI-powered pentest agent into HackerOne's top three U.S. business researchers while operating on an approximately $5,000 monthly budget and testing authorized production systems alongside experienced human researchers.

The three-month experiment tested whether agentic AI could deliver meaningful offensive security outcomes under real-world conditions rather than in a laboratory environment. FireCompass' system had to identify legitimate vulnerabilities, remain within the authorized scope of each program, provide proof of exploit, and withstand independent triage before its findings could receive credit.

During the April-to-June 2026 experiment, FireCompass recorded the following peak rankings:

- No. 3 on HackerOne's U.S. country board
- No. 2 for Highest Critical Reputation, ahead of XBOW
- No. 1 in the Up and Comers category

The positions were recorded as of July 17, following the completion of the primary testing period. HackerOne rankings are dynamic and recalculated daily, making the results a point-in-time measure of performance in a highly competitive environment.

The outcome provides a significant proof point for enterprises evaluating the role of agentic AI in offensive security. FireCompass had already achieved 100% results across the benchmarks it tested, but the HackerOne experiment imposed a much higher standard because invalid, duplicate, out-of-scope, or unverifiable submissions received no leaderboard credit.

The system submitted 150 reports during the primary April-to-June period and 204 across the full experiment. Nineteen reports were accepted as triaged or resolved, while another 58 represented genuine vulnerabilities that had already been reported by other researchers and were therefore classified as duplicates. Critical- and high-severity vulnerabilities accounted for 64.4% of findings with assigned severity ratings.

FireCompass conducted the experiment on an initial budget of approximately $5,000 per month, including AI tokens, cloud infrastructure, and human oversight. The company combined multiple frontier AI models with its purpose-built small language models and an automated pipeline designed to discover, validate, and report vulnerabilities.

The agents were deployed within defined operational safeguards throughout the exercise. These included hard enforcement of each program's authorized scope, non-destructive vulnerability validation, limits on request rates and concurrency, controls over potential blast radius, and proof of exploit for every submitted finding.

For EC-Council, the milestone provides tangible validation of an ecosystem strategy focused on combining advanced AI with skilled cybersecurity professionals, disciplined engineering, and accountable human oversight. It also supports the organization's view that the next phase of cybersecurity will be shaped by professionals who can direct, supervise, and apply autonomous systems effectively.

"FireCompass has taken agentic AI beyond benchmark performance and demonstrated that it can deliver credible security outcomes on live production systems," said Jay Bavisi, Group President, EC-Council. "This validates our belief that the future of cybersecurity is not AI replacing professionals, but AI-powered professionals operating with greater speed, reach, and precision. FireCompass is showing what becomes possible when advanced AI, disciplined engineering, and human oversight work together. For enterprises, this creates a clear path from periodic assurance to continuous proof."

The result also highlights a fundamental change in the economics of offensive security. Capabilities that once depended on large specialist teams, extended testing cycles, and significant budgets can increasingly be automated, scaled, and applied continuously across expanding attack surfaces.
That shift creates a substantial opportunity for defenders because security teams can test more frequently, validate exposures earlier, and focus human expertise on the risks that matter most. It also introduces urgency, as the same advances in AI capability and declining operating costs may become increasingly accessible to threat actors.

Bikash Barai, Founder and CEO of FireCompass, said: “Our agents had already reached 100% on the benchmarks we tested. This experiment was about what it costs to get into the Top 3 spots in a global bug bounty leaderboard. We reached there with just $5K per month. What it means is that threat actors with a very small investment can be as powerful as the world’s top hackers. This is an opportunity for defenders, but the same economics reach attackers at the same time. The next challenge is not just building more powerful AI. It is engineering the controls, safety, and accountability around it.”
FireCompass applies its agentic AI technology to autonomous penetration testing and red teaming across web applications, APIs, and infrastructure. Its platform discovers shadow assets and applications, validates which weaknesses are genuinely exploitable, and connects individual findings into multi-stage attack paths.

The company is trusted by Fortune 1000 organizations and has been recognized in more than 30 analyst reports. Its HackerOne performance strengthens the case for continuous security validation, where organizations move beyond periodic snapshots of risk and test their defenses as systems, applications, and attack surfaces change.

Methodology note: The rankings cited are peak positions recorded during and immediately following the April-to-June 2026 experiment. HackerOne rankings are dynamic and recalculated daily, and HackerOne did not sponsor, verify, approve, or endorse the FireCompass announcement.

About EC-Council
EC-Council is the creator of the Certified Ethical Hacker (CEH) program and a global leader in cybersecurity and AI education, training, and certification. Founded in 2001, the organization has certified over 400,000 professionals across 174 countries and is trusted by government agencies, defense organizations, and Fortune 100 companies worldwide. An ISO/IEC 17024 accredited certification body, EC-Council’s credentials are recognized by the U.S. Department of Defense under DoD Directive 8140 and are held by professionals across the Army, Navy, Air Force, and leading global enterprises. EC-Council operates across four divisions: Education, University, Services, and Technology. For more information, visit www.eccouncil.org.
Media Contact: press@eccouncil.org

About FireCompass
FireCompass is an Agentic AI platform for autonomous penetration testing and red teaming across Web, API and infrastructure. It discovers shadow assets and web applications, safely validates what is exploitable, and connects findings into multi-stage attack paths with near-zero false positives. Unlike traditional scanners, it discovers credential reuse, business-logic flaws, privilege escalation, and app-to-app or app-to-network lateral movement. It can operate autonomously or with expert-in-the-loop validation. FireCompass has 30+ analyst recognitions across Gartner, Forrester, IDC, and is trusted by Fortune 1000 enterprises. Try it free at firecompass.com/explorer.

Media Contact
FireCompass Media Relations

firecompass.com

EC-Council
Media Contact
press@eccouncil.org
Visit us on social media:
LinkedIn
Instagram
