hackerone

SANS Research Program

Product Briefing

AI with HackerOne:

Insights from the 2024 SANS Institute Survey

September 2024

AI is everywhere, and yet it requires thoughtful guidance and monitoring from humans to realize the promised benefits. The most effective approach is to deploy AI for the task at which it excels, while relying on human expertise to keep the door shut on malicious actors.

HackerOne

As powerful as AI is, it's not a cure-all for security risks. The best protection in a constantly changing threat landscape is a combination of AI processing power and human assessment skills. That starts, as the SANS AI Survey indicates, with training your people in AI basics, but it doesn't end there.

HackerOne backs up its AI threat detection with a defense force of ethical hackers who can fight back against malicious actors—even against attacks that count as "unknown unknowns," the ones AI can't spot.

We all want to stop vulnerabilities as early as possible in the software development process—on the left side of the flowchart, the way most people depict it. With HackerOne, you can start by integrating real-time testing for code changes and live applications, ensuring that any vulnerabilities are caught early and that code remains secure throughout the development cycle.

HackerOne also provides human-led audits of an organization's codebase and, as you might expect, penetration testing to make sure nothing has been missed. Unlike many pentesting organizations, a HackerOne test lets you, the client, monitor what the testers are discovering in real time and, if needed, shut down the test before the end date.

Key Findings

Most organizations are concerned about AI's impact on offensive cybersecurity tactics, with 79% worried about AI-powered phishing campaigns and 74% about automated vulnerability exploitation.

AI has limitations when it comes to detecting threats—70.5% of respondents indicated issues with AI generating false positives while missing outlier indicators or new threats.

Approximately 71% of organizations report higher satisfaction due to AI automating tedious tasks, allowing focus on rewarding work.

You can work with HackerOne to set up a limited-time challenge or a continuous bug bounty program, both of which deploy a global network of registered security researchers against your applications. In fact, some organizations use these challenges in the form of AI red teaming to test the robustness of their own AI systems, including their models and software components. If application security is mission-critical to your organization, this is the level of assurance you need (Figure 1).

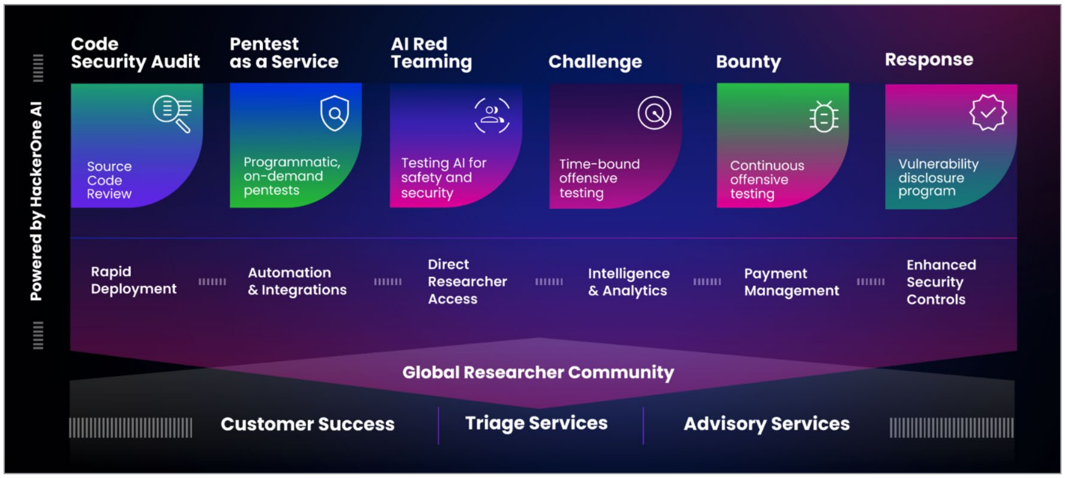

Figure 1. HackerOne Platform Capabilities and Comprehensive Offensive Security Testing Portfolio

Of course, HackerOne uses AI, for use cases like these:

- Tailored advice and personalized remediation guidance
- On-demand assistance with intricate reports, proofs of concept, and technical details
- Actionable insights from visuals and videos
- Custom apps for repetitive tasks
- Accelerated internal workflows and security tools with the HackerOne API

The latest product, Hai, is an AI co-pilot that adds context to vulnerability reports, transforms natural language into filtering queries, and uses vulnerability data from across the HackerOne platform to provide recommendations. The goal is to improve detection rates, prevent regressions, and speed up triage so analysts can spend more time fixing issues. Hai can read screenshots and the organization plans to introduce the ability to analyze videos that automatically detect key moments in videos, such as when payloads are sent to a system and executed. Integrating Hai into the workflows of people handling incoming trouble tickets helps front-line analysts call on the power of HackerOne's community through security-focused translation capabilities and auto-routing so that tickets quickly reach the team member best prepared to handle them.

The HackerOne team is clear that at least for now, it's not time to trust AI with making final decisions about security issues and mitigations. There's always a human pulling the trigger, and their goal is to make sure that humans are as well prepared as possible, no matter what comes in the door. That lines up with where a lot of the SANS AI Survey respondents fall on the subject of AI trustworthiness, with more than 40% saying they have difficulty understanding and trusting AI-based decisions due to a lack of transparency.

If your digital resources need to be protected no matter what, and if you believe the best way to stop a hacker is with a better hacker, HackerOne may be the solution you need.

For more information, visit
www.hackerone.com

Note that SANS Product Briefings do not represent a SANS endorsement of a sponsor or its products, but rather an overview of its offerings and their capabilities.