Securing the uncharted territories of AI systems. A discussion with Leo Racanelli

BY ELEANOR BARLOW · JUNE 11, 2026

The intersection of AI and cybersecurity is reshaping how we find, fix, and think about vulnerabilities. Yet for all the headlines, few conversations cut through the noise to ask what AI means for those on the ground: the hunters, the security engineers, and the organizations trying to secure their data.

In this blog, we open up that discussion, with insights from Leo Racaneli for an unflinching look at the state of play.

Leo Racanelli is a Senior Software Engineer, Security Engineer, Penetration Tester, and Bug Bounty Hunter from Italy. As a Bug Bounty Hunter, he's had the chance to report valid vulnerabilities to some of the world's largest companies across multiple industries, from tech to retail, automotive to media. His focus tends to be on flaws that require a bit of creativity to spot and are easy to overlook. He also takes part in live hacking events, where the real-time challenge adds a different kind of pressure, and is an Intigriti Ambassador.

What follows is an honest, ground-level assessment of where we stand today, and where we, as a world, are headed tomorrow.

Vulnerabilities seen in AI Applications

There is a big market hype on AI at present, and there are justified concerns when it comes to the security of AI. Nowadays, you can input AI into basically everything, not just web applications. But it does not necessarily mean that companies should. AI applications introduce a novel breed of risks. Understanding these unique exploits is the first critical step toward building robust, resilient architectures.

> **"Everyone feels that they need to have it everywhere. Before AI, you needed to know how to sanitize all the parameters. Of course, the protections are improving, but AI applications allow natural text input that can't be fully sanitized, which means that the data coming in and out isn't a fixed structure. There is a lot of rushing to market without thinking it through properly, without proper security knowledge, and without the due diligence to check that it will protect your software."**
>
> **Racanelli**

Companies need to get better at understanding what they control, but this is not an easy task, as the knowledge is not yet known.

These security considerations are not well understood even among experienced developers. In the race to market, many organizations overlook fundamental security guardrails, treating AI integrations as standard APIs rather than dynamic, unpredictable systems.

> **“Every software engineer with some experience knows what an SQL injection is, or an XSS. They know this even if they don't work in security because these things have been around for over 20+ years. But it's not the same with AI, and it takes some knowledge to understand how it can be manipulated.”**
>
> **Racanelli**

The appeal of bug bounty programs that incorporate AI

Traditional scanners often fall short against machine learning models, making human ingenuity more vital than ever. Incorporating AI into programs often incentivizes creative ethical hackers.

> **“For me, AI components in programs are inherently attractive to hunt. When AI is in a program itself, I want to try it as it usually motivates deeper investigation. For instance, when a program includes a chatbot scope, I always open all the references that I have in my notes for things like invisible prompt injection techniques from reference notes.”**
>
> **Racanelli**

How AI has impacted workflow and a pivot in focus

When thinking about how workflows have been impacted by AI, most will look at how integrating tools into security processes has reshaped operations. But for many, AI has impacted workflow in a very different fashion. It has become a working companion.

> **“I always wanted to collaborate more with other hackers, but it was difficult with timing for me. AI has become my teammate, my partner, in that I start asking it questions, and it comes up with something that I need to validate, and maybe I can give it some leads, some other time it gives me some leads, and for me, it's how it works the best.”**
>
> **Racanelli**

But some still question leaving AI itself alone.

> **“I know that many people do leave AI to run alone, and they are successful in that. I prefer to collaborate with the AI agent in funding new stuff, and it’s useful when I'm stuck, and I need a recommendation on what to do next.”**
>
> **Racanelli**

In a talk on ‘How AI will Transform Small Businesses’, former UK president Rishi Sunak stated that ‘The businesses making the most progress are not those with better technology; they are those prepared to adapt how their people operate around it.’ The same goes for how hackers are utilising the tools available to them.

The landscape of what ethical hackers actually target is also shifting. As security researchers pivot their focus, organizations must understand where these new digital battlegrounds are being drawn.

> **“Historically, my focus was all about authorization and authentication bugs. I really loved the applications with a lot of roles and users, which is still the case. But with the help of the AI, I was able to find even more bugs. The AI tries new attack vectors that I didn't try because I did not think about them. So, I'm reporting different bug classes, while still reporting the usual ones I would have searched for.”**
>
> **Racanelli**

The AI, in this example, occasionally identifies vulnerabilities, CVEs, or technology-specific issues that were previously unknown, functioning as a teaching tool.

> **“I learn from the Agent, but also the Agent learns from me, so it’s a knowledge exchange in both ways.”**
>
> **Racanelli**

Future-proofing decisions through a library of findings

By building a comprehensive library of AI findings, organizations can proactively identify systemic flaws before they manifest in future deployments. In fact, many hackers maintain a playbook, of sorts, that tracks historical actions to enable future actions.

> **“I maintain a playbook in Obsidian using markdown files with a tree structure. There is the Claude MD on top, with branching for specific targets and general knowledge. It tracks what has been found and what is not of interest to prevent the AI from repeatedly pursuing dead ends, and it documents both findings and non-bug leads to help free up context in future sessions.”**
>
> **Racanelli**

This knowledge can cross-reference similar bugs to identify differences and understand why certain approaches don't work.

> **“So, I might ask myself something like, "This is this bug, which is very similar to this other one. Can you check what the differences are and why it doesn't work?” It’s useful to find new bugs in other targets.”**
>
> **Racanelli**

The best and worst aspects of using AI in a workflow

Navigating the double-edged sword that is AI requires a careful balance of weighing up the pros and cons.

> **“The best thing, in my eyes anyway, as discussed above, is the companionship, the knowledge, the ability to find new classes of bugs. It teaches me and makes me a better hunter overall.”**
>
> **Racanelli**

But while having another teammate is arguably a good thing, there are some elements that are negatively impacting the crowdsourced community, especially when it comes to hindering new researchers who don’t have the tools in place to compete.

> **“In my eyes, the worst thing about AI is that it's a tool, and everybody is using it. The issue here is that the barrier to entry for bug bounty has increased due to this price wall, which is tokens. Bug bounty is no longer free to participate competitively; there is now a cost barrier.”**
>
> **Racanelli**

The key argument here is that you must be able to compete with others using the same advanced tools. For instance, live hacking events now require token budgets to remain competitive.

> **“If I want to attend a live hacking event, now I need to go there with all the tokens I can.”**
>
> **Racanelli**

In essence, AI is, arguably, gatekeeping the ability to compete in bug bounty.

> **“If you don't have AI, you can't be competitive. It's not optional, it’s essential, and it’s only getting worse. This is also why I'm focusing on not using all the tokens that I can, but rather focusing on the right usage of the AI. The goal is to learn that and spend the money that I need to stay in the game competitively, and not more.”**
>
> **Racanelli**

The next era of bug bounty programs. A new AI wave?

The next evolution of crowdsourced security may look a little different from what we know today. And the relationship between researchers and organizations may also change to match the speed of innovation.

> **“Current cheap token pricing won't last forever, so investing time now in proper usage is critical. My concern here is that if bug bounty costs increase, payouts may need to increase proportionally. So, you're investing time now to learn how to use it properly so that when token prices go up, you're not just wasting your tokens, which could be worth thousands.”**
>
> **Racanelli**

Right now, it is a fair exchange. But it’s only in the beginning stages. Maybe payout in the future will be in the form of tokens as well as currency. Only time will tell.

> **“Agentic AI has only existed for a few months, but it already feels essential, and technology is moving extremely fast with potential for significant changes in the next 3-12 months.”**
>
> **Racanelli**

The history of artificial intelligence is defined by waves of hype followed by periods of calibration and stabilization. Recognizing where we currently stand in this cycle helps security leaders distinguish between passing trends and lasting shifts.

Next steps

The conversation doesn't end with speculation; it ends with action.

- **Organizations**: Whether you are using AI to build faster or incorporating AI features into your products, this resource shows how to design an AI program suited to your security maturity. ‘AI Security & Safety’

- **Security Leaders**: check out '**CEO insights'** for opinions on pressing AI discussions, from our CEO and leadership team.
- **Triage team**: Interested in validation tools, read the AI impact from a **trager's perspective**.
- **Industry researcher**: Aware of all the AI misconceptions? **Test yourself here**.
- **Hackers just starting**: Check out **PortSwigger** and **OWASP** for more material.

Have any questions on the topics covered in this content? Please **contact the team.**

AUTHOR

Eleanor Barlow

Eleanor Barlow is a London-based Senior Cyber Security Technical Writer at Intigriti, with 9+ years' experience reporting on and writing for the cyber and tech sector. She specializes in data-driven content on cybersecurity and bug bounty intelligence, helping organizations benefit from the latest trends and insights.