



Using AI the smart way. Interview with Cristian Zot (CristiVlad25)

BY ELEANOR BARLOW · JUNE 17, 2026

Cristian Zot, known by most in the industry as [CristiVlad25](#), is an active security researcher, experienced pentester, and an Intigriti [Hacker Ambassador](#).

He is a prominent figure in the ethical hacking community and frequently collaborates with Intigriti through platform meetups, podcast appearances, and educational content. Cristian has featured as a guest expert on Intigriti's live Office Hours podcast session on Discord, taking community questions.

More recently, he served as a guest speaker at Intigriti's Bounty Sync event in London, leading discussions on AI.

Today, we continue that conversation and delve into the impact AI is having on pentesting and Cristian's insights on how to use AI the smart way.

How useful do you find AI, and how much do you use it daily?

I use AI heavily in my day-to-day work. A lot of what I share cuts against the hype you often see online, where the fantasy of fully autonomous hacking bots that find bugs while you sleep is abundant.

The reality is that AI is far more useful and a lot more grounded than that. Think of it as a force multiplier, but only when you stay in the loop.

You need to know how to fold AI into your workflow without handing over your judgment, why "autonomous" is still a myth, and where your value as a hacker comes from.

Great, let's dive in! What kind of experience do you think is important to get started?

Ideally, you will have had a few months of experience under your belt, have landed some valid bugs, collected a few dupes, made the occasional [out-of-scope](#) (OOS) mistake, delved through a pile of [PortSwigger labs](#), and want to start using AI but aren't sure how to do it well.

If that sounds like you, the goal here isn't to hand the work over. It's to move faster and learn deeper while keeping your hands firmly on the wheel.

What would you say is the first step?

Before anything else, give the AI your scope. This keeps suggestions relevant and, just as importantly, helps you avoid the OOS mistakes that burn so many new hunters. If the model knows what's in and what's out of bounds, it won't send you chasing assets that'll get your report closed as out of scope.

Pick your mode based on what you're trying to do. AI isn't one tool; it's several, depending on what you ask of it.

If you are going after a particular vulnerability type, then frame the conversation around that class. Tell the AI what you're looking for and let it help you reason about where and how it might show up in your target.

If you are going deeper or verifying something, when you want to dig into a behavior or confirm a finding, share your requests and responses with enough context for the AI to follow what's happening. Ask for guidance on where to test next, rather than a simple yes/no verdict.

What about when it comes to learning?

This is where AI quietly shines. Ask it three things:

- Ask it to explain concepts visually
- Ask about error codes you don't recognize
- Ask for a visual breakdown of something like SSRF

Feed Request for Comments (RFC) into NotebookLM and let it generate a breakdown. It will even produce an audio walkthrough you can listen to on the go or a video if you're a visual learner.

What about improving reports? How can it benefit researchers there?

Write the report yourself, and then, after you've verified the bug, only then bring in AI to clean it up. Never get AI to write it for you.

My top tip is not to let AI inflate your report with paragraphs of filler. Tell it to be concise, with bullet points instead of paragraphs, because triagers want signal.

Read this guide on using [AI for improved vulnerability report writing](#) before you get going.

What is your advice to those wanting to wire AI into everything?

It's tempting to do this, but it is important to resist this. Don't go from 0 to 100 overnight. Introduce AI gradually so you understand what it's adding and, more crucially, what it's getting wrong.

It's important to be able to understand the limitations. Where the model breaks down is what stops you from trusting a confident-sounding mistake.

What do you believe to be the biggest pitfall with AI?

AI is there to please you. This one's worth sitting with. Models are optimized to be agreeable. They'll validate a shaky theory, "confirm" a bug that isn't real, and generally tell you what you want to hear. In security work, that tendency is dangerous because false positives waste everyone's time.

I recommend that you treat enthusiastic agreement from an AI as a prompt to verify, not as confirmation.

Are there any common misconceptions you would like to debunk?

There's a popular idea floating around that we already have autonomous pentesting. Point a tool at a target, walk away, and come back to a pile of bugs. I don't think it's that simple.

In practice, the human should still be in the loop. You provide the input, you review the output, you answer the model's questions, and you make the calls. AI accelerates the work; it doesn't own it. Anyone selling fully "autonomous" hacking may be disingenuous.

If AI handles the mechanical work, where does the value come from?

Here's the hard truth: if you don't have a skillset, you can't add value. AI amplifies what you already know; it can't substitute for understanding you don't have. The hunters who get the most out of these tools are the ones who can spot when the output is wrong and steer it somewhere better.

So, that means you need to keep building real depth. The more you know, the more the AI is worth to you, not less.

You are the validation layer. Before anything goes into a report, you confirm it's real, reproducible, and in scope. The AI proposes; you verify. Every time.

Models only work with what you give them. As you learn more about a target, new endpoints, auth flows, and odd behaviors, keep feeding that context back in. The quality of what you get out tracks closely with the quality and freshness of what you put in.

The [Intigriti Hackademy](#) is a solid place to keep leveling up.

What is your key takeaway for those looking to use AI tools in bug bounty?

AI in bug bounty isn't a replacement for skill, instinct, or judgment; it's a multiplier on all three. Start small, stay skeptical, keep learning, and never let a model's eagerness to please stand in for your own verification.

The hackers winning with AI aren't the ones who handed off the work. They're the ones who got better at directing it.

Next steps with Intigriti

A big thanks to Cristi for the interview. Now go put his advice to the test!

Check out our [public programs](#) and try out your AI-assisted workflow on a live target!

Or, if you would like to know more about how to use AI, visit our [AI Security and Safety](#) page for further details.



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.

REQUEST A DEMO

intigriti.com/demo

VISIT THE WEBSITE

intigriti.com

GET IN TOUCH

hello@intigriti.com