


Inside the Mind of a **Hacker**

The Human-Augmented Intelligence Era



How would you define “Hacker”?

Today 9:41

If you were to ask 10 people off the street whether they could distinguish between hackers and cybercriminals, they'd likely be unable to do so.

Merriam-Webster defines a “hacker” as “an expert at programming and solving problems with a computer.” While “hacker” is the predominant self-descriptor used by the cybersecurity community (with even some CISOs we know adopting the moniker), this benevolent term has sadly become synonymous with malice. The bad guys also call themselves hackers, and unfortunately, they get most of the attention.

Here at Bugcrowd (and in this report), we refer to the good guys as hackers. Other terms you may have heard include “ethical hackers,” “white hat hackers,” and “security researchers.”

expert

ethical

security researchers




Table of contents

LETTER FROM THE CEO

4

CHAPTER 1

Hacker Demographics and Motivations

INFOGRAPHIC	Double Click on Hackers	6
ARTICLE	Meet the Hackers Behind the Screens	7
SPOTLIGHT	The HackHer Network	11
ARTICLE	Understanding the Complexities of Hacker Motivations	14

CHAPTER 2

Hacking in Teams

ARTICLE	The Power of Hacking in Teams	19
SPOTLIGHT	Meet an Elite Hacking Team	22

CHAPTER 3

Hackers and AI

INFOGRAPHIC	3 Years of AI Innovation	26
ARTICLE	The AI Hacking Advantage	27
SPOTLIGHT	Meet CISO and Hacker Aaron Guzman	31
THOUGHT PIECE	Top Ways Hackers are Using AI	34

CONCLUSION

38

Letter from the CEO

Every organization in the world is inherently vulnerable. On top of that, the attack surface is expanding, and humans can't scale it alone. As a result, many have pointed to the "AI revolution" as a way to approach an impossible situation for security teams.

So much of the conversation around AI has been a tug of war. Are you Team Humans or Team AI? The discourse seems to say, "Pick one, because only one will come out the victor."

However, I just don't see it that way. **I'd argue that we're actually entering the human-augmented intelligence era:** an inflection point where human ingenuity and the power of AI are converging to make each other better.

This report examines this convergent approach. It uncovers the human side—the hackers behind the screens, their unique skill sets, their motivations, and their shift toward teamwork. It also looks at the AI side of the equation, breaking down how hackers are using AI as a force multiplier, expanding their impact exponentially.

I was particularly struck by a quote from "**Top Ways Hackers are Using AI**" (page 34), when OxMoose states, "Security is a race where the finish line keeps moving forward and the only way to stay ahead is to run faster than everyone else."

By utilizing AI, the pack not only gets 'smarter' but also quicker, running on a relentless 24/7 schedule."

In this edition of *Inside the Mind of a Hacker*, we interviewed and surveyed over 2,000 of the world's leading hackers. The findings from this research can help organizations better understand the hacker community while leveraging the scale that machines and AI bring to build more resilient security programs.

Ultimately, this is why we at Bugcrowd are passionate about realizing an intelligent, self-learning Platform that unites human creativity with machine learning, giving both customers and hackers endless opportunities to make smarter, faster decisions that help preempt emerging threats. And we're just getting started! ■



David Gerry

Report Highlights

This edition of *Inside the Mind of a Hacker* analyzed over 2,000 survey responses from hackers on the Bugcrowd Platform, in addition to hacker interviews.

WHO THEY ARE

92%

are 34 years old or younger.

98%

are proud of the work they do as a hacker.

56%

believe hacking is becoming more about geopolitics than curiosity.

81%

of part-time hackers work in a security-related field.

69%

graduated from college.

20%

identify as being neurodivergent.

HOW THEY ARE WORKING

72%

believe that hacking in teams yields better results.

95%

believe hacking is an art form.

74%

believe AI has increased the value of hacking.

82%

use AI as part of their hacking workflow.

61%

find more critical vulnerabilities when hacking in teams.

WHAT THEY ARE DISCOVERING

65%

have chosen not to disclose a vulnerability because there was no clear pathway to report it.

85%

believe reporting a critical vulnerability is more important than making money from it.

71%

have found a new vulnerability in the past 12 months.



Double click on hackers

88%
are 18-34

79% speak two or more languages




19%
identify as neurodivergent

HACK AS WORK

- 66%** hack part time
- 48%** say hacking helped them get a job
- 71%** encountered a new vulnerability type in the past year



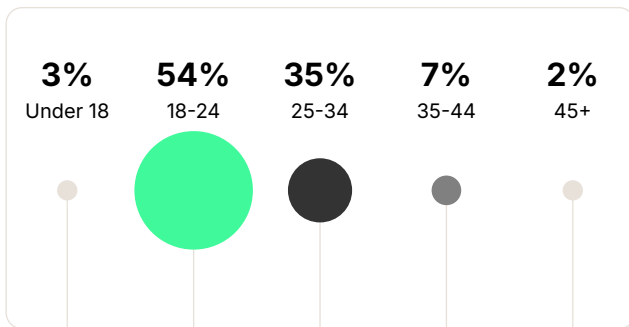
61% have hacked for the "lulz"




Meet the Hackers Behind the Screens

The Basics


Age



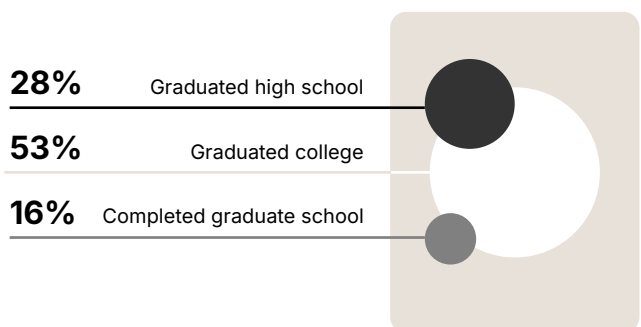
Gender



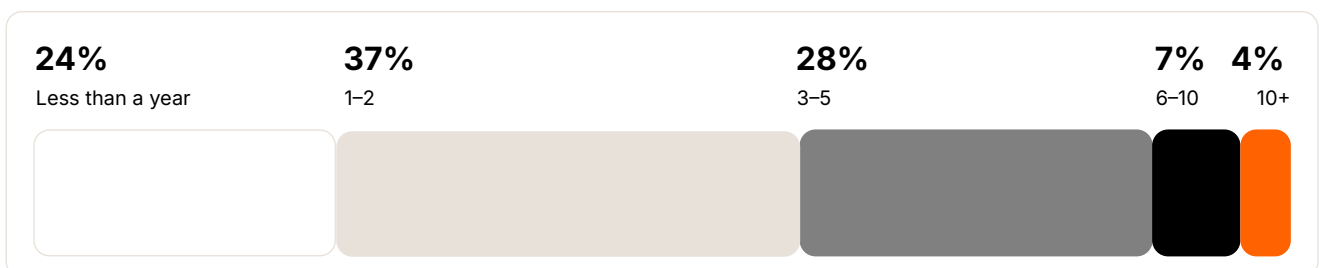
Languages spoken



Level of education completed



Years hacking



Workin' 9-5

Most hackers primarily work as hackers or IT and security experts, but they also work in lots of other roles, from architecture to the arts.

Top 10 primary occupations

Hacking and security research	01
Information and web application security	02
Education, training, and libraries	03
Architecture and engineering	04
Office and administrative support	05
Business and financial operations	06
Management	07
Installation, maintenance, and repairs	08
Marketing and communications	09
Arts, design, entertainment, sports, and media	10

Where in the world are the hackers?

Top 10 Countries

-  **India**
-  **Bangladesh**
-  **Egypt**
-  **USA**
-  **Pakistan**
-  **Nigeria**
-  **Nepal**
-  **Kenya**
-  **Indonesia**
-  **Turkey**

Neurodiversity in hacking

1 in 5 hackers identify as neurodivergent



The hacker community has quietly become a haven for neurodiverse individuals, with 1 in 5 identifying as neurodivergent. Hacking demands deep pattern recognition and systems thinking, skills that often align with neurodivergent strengths. The job also requires empathy—the ability to understand people and processes at a fundamental level—which tends to create a more tolerant and inclusive culture.

Furthermore, the masking strategies that neurodivergent hackers develop to navigate everyday social situations can translate into valuable professional skills. Social engineering, for instance, requires reading people and shaping interactions to achieve specific outcomes.

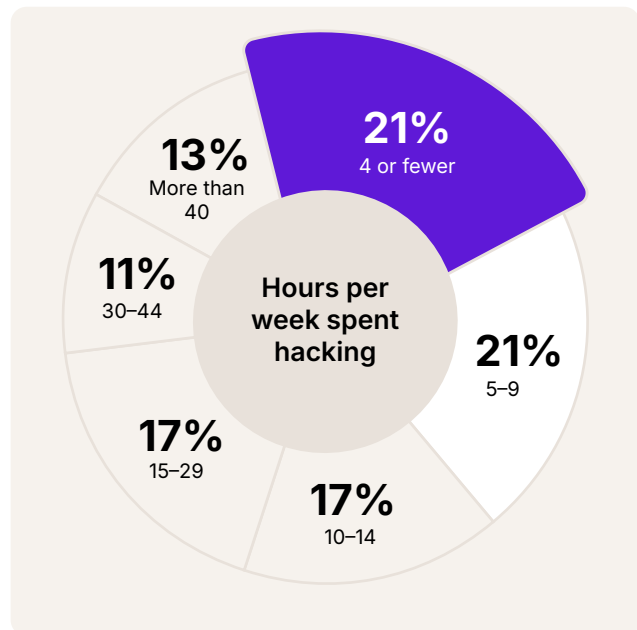
The same tools one uses to "pass" in neurotypical spaces become assets in understanding human behavior.

Skills to pay the bills

Nearly half spend fewer than 14 hours a week hacking, with over 40% keeping it under 10 hours. Part-time hacking is often offset by full-time security roles. When CISOs engage with the Crowd, they're tapping into a pool of security talent and expertise that many don't have the budget to hire full time.

Hackers are constantly building their skill sets to address new threats, stay sharp, and compete in the talent pool.

Security teams benefit from this, getting access to emerging skills and the diversity to augment their internal teams.



How hackers stay sharp

Practice hacking on new targets **80%**

Study tutorials/blogs/videos **75%**

Participate in CTFs **59%**

Take courses **35%**

Invest in certifications **31%**

Attend events/LHEs **28%**

Connect with mentor **25%**

Hackers' strongest skill sets

Network and infrastructure security **64%**

Software and application security **62%**

Cloud security **30%**

Mobile security **28%**

Social engineering **24%**

Most valuable skill sets

Software and application security **53%**

AI **51%**

Cloud security **50%**

Network and infrastructure security **43%**

Mobile security **37%**



85%
believe the Bugcrowd Platform offers more opportunities than other platforms.

What's in a (hacker's) name?

Hackers resist being boxed in. 59% call themselves either security researchers or hackers, closely followed by pen testers (57%).

But when forced to pick just one label, the split is telling: 28% lean into security researcher, 24% claim hacker, and 18% go with pentester. The title might change depending on the room, but the work stays the same.

How hackers identify

Security researcher	59%
Hacker	59%
Pentester	57%
CTF player	33%
Red teamer	33%
Security professional	26%
Security engineer	23%

Hackers' primary identifier

Security researcher	28%
Hacker	24%
Pentester	18%
Red teamer	9%
Security professional	8%
Security engineer	6%

The HackHer Network

The first women-only hacking community

There's an untapped talent pool—women in cybersecurity. Unfortunately, the cybersecurity industry has struggled to create inclusive pathways that welcome and retain diverse talent. As a result, only 5% of hackers identify as women, and this percentage has remained relatively unchanged over the past few years.

At Bugcrowd, we want to do our part in creating a hacking community where everyone thrives. That's why we launched the HackHer Network, a community for women in cybersecurity built by women who want to connect and grow together.

How the HackHer Network works

The community is hosted in a **Discord group**, where members can connect with one another and access exclusive community resources. They can opt for mentorship opportunities, technical challenges and skill-building events, networking, business partnerships, and content and media opportunities.

Only 5% of hackers identify as women

An overview of the HackHer Network

The goal of the HackHer Network is to create a space where all women, including cis, trans, and nonbinary individuals, can connect, collaborate, and grow together in a safe and productive online environment. Our community welcomes women at all career stages, from those just starting out to hobbyists and seasoned professionals.

By connecting inspiring women who have navigated similar challenges, we're building a network that empowers current members of the cybersecurity community while inspiring the next generation of hackers and cybersecurity leaders. Since launching in March 2025, our community has grown to 200 members spanning 30 countries, with strong representation from historically underrepresented regions such as Nigeria, Kenya, Uganda, and Ghana.

Member spotlights

Organizations need a steady supply of fresh perspectives and new voices to stay ahead of emerging cybersecurity threats. Meet three of our members whose stories capture the spirit of HackHer and show what it means to thrive together.

Katie Paxton Fear

Also known as [InsiderPhD](#)

Katie's extensive professional background ranges from game programming to software development. After deciding to pursue a PhD program, she found herself studying cybersecurity. She discovered that her background as a software engineer provided her with unique advantages. "I really enjoyed API hacking because I used to make APIs, and I realized that I can very quickly understand how they were built and how to break them," she explains.



YouTube @InsiderPhD →

Community and mentorship have been a huge part of Katie's success, and that's why she's giving back as a content creator and an educator. She [create educational content](#) that provides thorough yet accessible information on bug bounty hunting. Her library of videos constitutes the course she wishes she had access to before she pursued this path.

Her authenticity is also inspiring. She says, "I've always demanded a space for myself. This is my interest, my field, my career, and I'm going to own it."

“””” Given how competitive bug bounty hunting is, you'd expect us to all be shielding our computer screens and hissing at you, but actually, I've found that the majority of top hackers really want to share—they're as passionate about hacking as you are!

Brigitte Lewis


Also known as (t00t_toot)

“Women get paid less, experience more discrimination, and leave the industry in really high numbers, which is reflective of just how difficult it can be to thrive in an overwhelmingly male environment.”

Brigitte Lewis is a former sociology lecturer turned ethical hacker based in Melbourne, Australia. With seven years of experience in bug hunting and a PhD in sociology, she has become a vocal advocate for women in cybersecurity while building a successful career as a pentester and security consultant.

She credits the mentorship she received from women as critical to her success and passion for hacking, which continues to this day.

However, her career success has come with a sobering realization about the industry's gender problems. **That motivated her to support other women by starting W0m3nWh0HackM3lbourn3,** a space for women to hack and learn together while having fun.



[W0m3nWh0HackM3lbourn3](#) →

Olufela Osideko


Hailing from Ibadan, Nigeria, Olufela found the cybersecurity field through Africa's first women-only cybersecurity program.

She quickly found her groove with offensive security, particularly in penetration testing. It felt like her calling, something that she's "obsessed with."

Now, she channels that experience into advice that she shares with other women in the network. She says, "I wish I had prioritized myself more. I also wish that I had ignored the people who discouraged me....

Never doubt yourself. The world is waiting for you, and we need you!" ■

“Hacking gives me a sense of purpose. There are a couple of things I have tried in my life, but nothing comes close to hacking. I am obsessed with it; if I had all the time in the world, I would continue learning about it as much as I can.”




Understanding the Complexities of Hacker Motivations

Hackers are often portrayed in extremes and contradictions—either as a force for good or as the minions of villains—but this binary framing oversimplifies human ethics. The evolution of hacking itself reflects this **PARADOX**.


Hacker culture was born in underground, fringe spaces like basements, chat rooms, and bulletin boards. Today, corporations and governments actively recruit from that same culture, making what was once subversive part of the establishment.

Hackers have complex, nuanced, and evolving motivations that are driven by a push-pull of idealism, curiosity, and self-interest. The lack of a simple answer might seem like a drawback, but it's actually the key to recruiting and retaining hackers to join your side.

How do you think the general public perceives you as a hacker?



How do you think companies perceive you as a hacker?



Our survey results aren't surprising. Companies have a long history of working with hackers through crowdsourced security programs, like bug bounties, red teaming, or pen testing. The general public doesn't have this context, so they only know what they see in the news or movies, which can be a mixed bag. This difference in perception isn't bad—it's simply a reflection of the messy nature of hacking, which also spills over into how hackers think about their work.

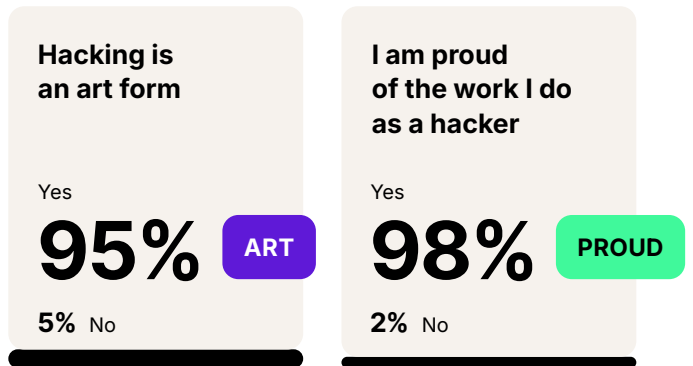
So, why do hackers hack?

When we ask hackers why they hack, we get a lot of different answers.

What words best describe why you hack today?

Financial gain	74%
Opportunities	52%
New experiences	51%
For the greater good	36%
To break things	30%
For recognition	29%
For the thrill	26%
Hacktivism	21%
Power	10%
Patriotism	6.4%
Clout/street cred	5%
To get people to do what I want	4%

Based on this list, it might be easy to conclude that most hackers hack for financial gain and that hacking is just a job to many. But in reality, multiple things can be true simultaneously. Hackers are just like everyone else; a paycheck is one of the primary reasons why they show up to work. However, most hackers view hacking as an art form, not just a way to support themselves. They take pride in their work and find purpose in it.



These beliefs show up in how hackers prioritize various demands. For instance, if a hacker identifies a critical vulnerability, their drive to help usually outweighs their desire to make money. However, if there isn't a clear or trusted channel to report an issue, they might avoid reporting it out of fear of being prosecuted, blamed, or ignored.

Reporting a critical vulnerability is more important than trying to make money from it



What happens to these unreported vulnerabilities?

Some stay lost, others are found internally, and unfortunately, attackers often find these entry points. It is crucial for security teams to improve their processes and provide a clear way for hackers to safely disclose their findings. A great way to do this is through a VDP.

I have chosen not to disclose a vulnerability because the company lacked a clear pathway for me to report it



Taking a page from spies


Hackers aren't unique in having complex motivations. Psychologists studying espionage identified a similar pattern and mapped out four core motivators for espionage: money, ideology, coercion, and ego (i.e., the MICE framework).

At first glance, it might seem like we can categorize all spies into one of four camps, but this would be an oversimplification.

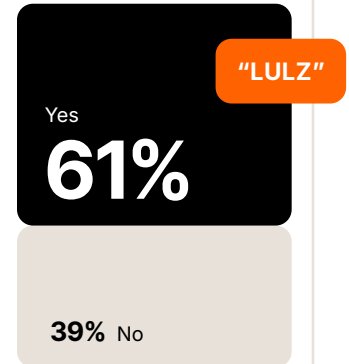
First, motivations often overlap. An individual driven by ideology might also be motivated by ego, and this overlap is actually advantageous—people motivated by multiple factors are more likely to help your cause.

Second, motivations like ego and coercion might carry negative connotations, but they really represent a human desire to be seen and to protect ourselves and our loved ones. These aren't inherently good or bad things; they're just human needs.

Hacking is becoming more about geopolitics than curiosity or fun




When hacking, have you ever done something purely for the "lulz"/laughs?



When we apply this framework to hackers, we see something striking: their motivations are increasingly being shaped by external forces. As economic inequality rises and opportunities shrink, there's more pressure on young people to find and create economic opportunity. Escalating geopolitical tensions are driving nation-states toward cyberwarfare, causing them to recruit hackers for state-sponsored operations. Hackers themselves are feeling this pressure, with many reporting how hacking is slowly becoming more about money and geopolitics than curiosity or fun. But that doesn't mean that their desire for curiosity and fun has disappeared. In fact, many hackers still pursue a target purely "for the lulz" (i.e., out of a playful, mischievous spirit).

This passion and curiosity needs to be harnessed and cultivated, not ignored.

Hacking is becoming more about money than curiosity or fun



How to tailor your program to appeal to hacker motivations

Connecting with hackers requires acknowledging that their motivations are complex, layered, and often contradictory. This sentiment needs to permeate every aspect of how we work with hackers, from messaging to program design, incentive structures, and recognition systems.

Let's break down what this might look like in practice. **Knowing that the top three hacker motivations are financial gain, opportunities, and new experiences,** security teams can design their crowdsourced security programs as follows:

- Competitive rewards** Ensure your reward ranges are competitive to attract elite hacker talent to your program.
- Broad scope** To satiate hacker curiosity and create growth opportunities, offer a wide scope across unique targets to pique hacker interest.
- Safe harbor** Hackers do their best work when they don't fear legal repercussions. Establish a policy that allows hackers acting in good faith, as defined within the program, to provide security feedback.

It's time we stopped trying to make ethics simple and started making our arguments sophisticated enough to match the reality of the people we're talking to. ■

A CISO + hacker's advice

How to maximize your investment in working with hackers


Organizations that are serious about security invest in scaling their capabilities. Crowdsourced security methods are key levers—not just for finding bugs but for building communities around your products.

The hackers who know your systems best become advocates and force multipliers for your security program. Embrace hackers as an extension of your team, not as vendors or adversaries you're reluctantly cooperating with. As partners focused on exposing quality defects and breakdowns in your processes, researchers don't offer findings as criticisms—they're gifts that make your products better.

Here are my 5 top tips →

Aaron Guzman

CISO, Cisco Network Product Engineering



1 Be responsive

Submissions disappearing into a black hole kills engagement. Even a quick acknowledgment matters. Hackers talk; responsiveness lends itself to your reputation in the community.

2 Define scope thoughtfully

Overly restrictive scope signals that you're checking a box and not serious about security. If large portions of your attack surface are off limits, examine why. Sometimes, legitimate reasons exist, but often, these limits stem from organizational fear worth addressing. Your scope evolves as risk indicators evolve, so ensure there's reciprocation.

3 Pay fairly

Bounty amounts signal how you value security. If critical vulnerabilities pay less than a consultant's day rate, you're telling hackers that their work isn't valued. They'll invest their time elsewhere.

4 Build relationships

The best results come from hackers who understand your environment deeply. Develop ongoing relationships with those who demonstrate quality work. Give them more access over time. Some of my most valuable security insights came from hackers who've tested our products for years. You might even end up hiring some onto your team as pentesters or security engineers.

5 Close the loop

Tell hackers when and how issues were fixed. Share what you learned. This builds trust and often leads to follow-up research that reveals the next layer of issues.

The Power of Hacking in Teams



Real-world threat actors understand the power of teamwork. Bob Lord, former Chief Security Officer at the DNC, described what his team faced from nation-state hackers: **"We're up against dedicated human adversaries. They work in teams. If one team is failing in some way, they get to call in another team."**

State-sponsored attackers organize into specialized units with distinct expertise, and ransomware operations involve multiple groups working together. Initial access brokers find entry points, ransomware-as-a-service operators develop the tools, and execution teams deploy the final payload. Different people handle different aspects of the attack chain, each bringing their specific skills to the overall operation.

Think Ocean's 11

You need a diverse crew made up of individuals who excel at different things to pull off a heist. When you connect those skills together, the whole becomes greater than the sum of its parts.

Security researchers are adopting the same collaborative approach, and the results speak for themselves.

Surveyed hackers report that 40% of them currently hack as part of a team while another 44% want to hack with a team but haven't found teammates. When hackers work in teams, they can improve their skills by learning from each other, combine strengths, expand professional networks, and even earn higher pay. Nearly half of hackers, around 44%, report making more money through team collaboration.

You may already be familiar with hackers working in teams. Red teams simulate real-world attacks against an organization's defenses.

Breaches require reconnaissance, initial access, lateral movement, and privilege escalation. Just as attackers organize into specialized units, red teams need diverse skill sets working in coordination.



Why teams matter more than ever

The case for team-based hacking is also about keeping up with tech itself. The skill set needed for security testing resets approximately every two years. What worked in the web era changed with mobile, then cloud, and now, AI, which is accelerating the cycle even faster. Each wave brings new attack surfaces, architectures, and vulnerabilities that require different areas of expertise.

No single person can constantly relearn entire skill sets that quickly.

The technology stack has become so abstracted and layered that mastery across all areas is pretty much impossible. Even the best hackers have blind spots, and these are not due to a lack of skill or dedication.

Together, teams stay ahead of the curve and provide coverage that a single individual cannot. The attack surface keeps expanding, but teams expand with it.

How companies benefit from hackers working in teams

So what do you actually get from team-based testing?

→ **Cost-effective coverage across your entire attack surface**

72% of hackers agree that organizations get better results from teams than from individuals. Organizations benefit from the multiplication effect of diverse expertise applied simultaneously.

→ **Faster results**

About 80% of hackers agree that team hacking is faster than working solo. While one hacker might spend three weeks thoroughly assessing a web application, a skilled team can complete the same work in days through parallel processing.

→ **Higher-quality results through combined expertise**

No single hacker masters every technology stack or attack vector, but teams come closer to doing so. When asked directly, 60% of hackers find more critical vulnerabilities when they hack as a team.

What makes teams effective

Not all hacking teams perform equally. The most successful teams have these traits in common:

A track record of trust and communication

Team members understand each other's strengths, delegate accordingly, and move faster because there's no need to constantly double-check work or second-guess decisions.

Balanced skill levels and complementary specializations

The most effective teams consist of people at roughly similar skill levels who bring different specializations.

Small size for coordination

When asked about ideal team size, 70% of hackers say 3–4 members. At this size, a team has enough diversity for complementary skills while remaining tight enough to foster real relationships and uphold high accountability standards.




Matching your defense to reality

Today's attackers often work together in teams, devising complex, chained exploits and creative attacks that require the expertise of a team to mitigate.

The future of effective security research lies with teams. If attackers are working in teams, defenders should do the same. ■

Meet an **Elite** Hacking Team

The difference between good teamwork and elite teamwork isn't about adding more people—but the right ones.



One team has turned this principle into a winning formula. Meet sw33tLie (Paolo), bsysop (Guillermo), and godiego (Diego): three hackers who work together as an elite hacking team.

Working across different time zones and languages, these three have developed a collaborative approach to quickly identify critical vulnerabilities, which recently earned them top

honors at Bugcrowd's Hacker Showdown. But what makes their partnership work so well?

In this article, we'll hear directly from the trio about how they formed their partnership, their secret sauce to working well together asynchronously, and what organizations can do to attract elite teams like theirs.

Hacker teams work in different ways! While some teams focus on hacking together at live events, others prefer doing it on a daily basis for regular hacking on security programs.

 sw33tLie, bsysop, godiego



 Messages  Files +



Coming together as a team

Sw33tLie, bsysop, and godiego were individual bug bounty hunters before they formed their team. Bsysop explains how they met,

“I met the others [godiego and sw33tLie] over Slack. We started by talking about random things and sharing tips, which built rapport, before meeting in-person at a live hacking event in Tokyo,” bsysop explained.




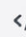
This focus on friendship first enabled each team member to get to know the other in a low-stakes environment, organically building rapport and an understanding of each other’s values. “One thing that’s really important to all of us is the value of personal relationships and enjoying the process,” Godiego explained. “Small things like sharing memes and talking about random stuff are really important to us.”

It’s also remarkable because all three members represent different generations of hackers. Sw33tLie is the youngest of the group, followed by godiego, and then bsysop. But this intergenerational setup actually enhances their work, with more experienced team members offering mentorship while younger hackers contribute fresh approaches and energy. Sw33tLie notes,






“Age is just a number in this field. What matters is curiosity, skill, and the willingness to learn from each other.”



As a result, they built enough trust to take the leap and work together, starting with smaller collaborations before scaling up. Today, the team primarily focuses on hacking together at live and other special events.



B I U    

Me Private invites at 2am

 **Aa**    

Shit + Return to add a new line



A breakdown of their process

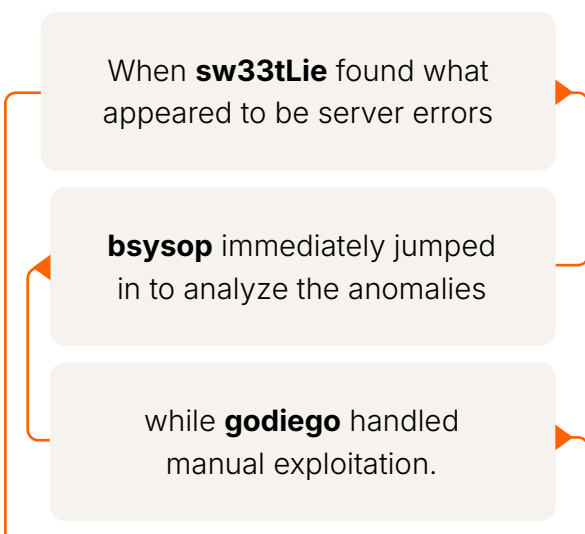
Instead of adhering to rigid roles, the team adopts a flexible approach in which everyone constantly shifts between leading and following as needed.

“Some teams have a lot of big egos and try to be rock stars, but we don’t play like that,” bsysop explained.

“We’ve kept a no-ego culture: if someone finds a better way, we adapt fast. That mindset lets us scale without losing the creative edge that got us here.” This flexibility also extends to how they continuously improve their collaboration process—after every project, they conduct a retrospective to identify changes to work even more effectively together.

This ego-free foundation allows each member’s strengths to shine. “Our strengths complement each other, like how one person is good at exploit development while another person is stronger at creative recon and changing logic. Together, that mix lets us see the full attack surface, execute faster, and bring forth creative ideas and techniques that lead to interesting exploits,” Godiego explained.

This creativity was exemplified during the Hacker Showdown, where their complementary strengths enabled them to discover a rare function injection vulnerability.



Working together, they realized these weren’t glitches at all—they’d uncovered a previously unknown vulnerability. With no documentation available, they reverse-engineered the exploit through collaborative experimentation, ultimately finding a way to bypass function errors and access sensitive data, securing their victory.



Working asynchronously

Working across different time zones is usually challenging for most teams, but the trio has turned this potential obstacle into an advantage. Their solution is an “organized chaos” approach that relies on constant asynchronous communication through a single Telegram thread, ensuring someone is always making progress regardless of who’s awake.

While this approach can get messy at times, their shared history means they’ve developed enough shorthand and context that it’s easier to use Telegram chats than formal, organized tools like Slack channels. They deliberately choose their cluttered Telegram thread over structured Slack channels because, as bsysop puts it,

“We know each other, I know exactly the mess sw33tLie is creating. So it’s our mess.”

Perhaps most importantly, they’ve learned that maintaining a sense of fun and experimentation is crucial for surviving the stress and long hours across time zones. “When you’re working together, you can’t focus on the money. You have to remember why you’re there: you’re all just crazy guys trying to hack and find new ways to break something,” Godiego noted.

That playful mindset not only prevents burnout but also creates opportunity for new discoveries. For instance, when sw33tLie and bsysop became curious about HTTP smuggling techniques, they chased their curiosity with another hacker, leading to the discovery of a [critical vulnerability](#) across thousands of Google Cloud websites.



Why (and how) organizations should embrace hacking teams

The trio's success highlights the unique value of teams in crowdsourced security programs: their ability to discover vulnerabilities that solo hackers might miss entirely. "When working alone, you might find a lead and investigate it, but if you don't find anything, you'll move on," Sw33tLie explained.

"When working as a team, we can bring in different skillsets so it's easier to keep digging, which leads us to discover more novel vulnerabilities."

Another side benefit: teams provide better program support through faster response times. With members spread across time zones, someone is always available to communicate with program owners.

Given these advantages, **what can organizations do to attract more teams to their crowdsourced security programs?**

The trio offers four practical tips:

Click here and reveal the memes

i-am-normal

being-a-nosy-nerd

1 Allow team submissions

Explicitly welcome team participation in your program guidelines. It's a simple step that encourages team participation.

2 Reward collaboration

Incentivize teamwork by adding shared bounties to your rewards and by recognizing teams on leaderboards.

3 Showcase team success

Publish blog posts of successful team discoveries to signal that you value collaborative contributions.

4 Make your scope flexible


Teams often bring custom automation or tooling that can uncover deep issues. An open scope gives teams the space to leverage that knowledge and explore, resulting in the discovery of higher-impact vulnerabilities.

Teamwork makes the dreamwork in all areas of life—including crowdsourced security. By embracing teams, organizations can access more expertise and creativity to stay ahead of sophisticated threats. ■

3 Years of AI Innovation


We asked hackers these four questions about their generative AI usage and beliefs over the past three years.

Check out how their responses have evolved!



Hackers using generative AI for hacking

More hackers are adopting generative AI technologies in their security research workflows.




Top use case for generative AI in hacking

From automation and speeding up operations to improving the accuracy of hacking tools and analyzing data, hackers are finding new ways to leverage AI.




Hackers who believe AI technologies outperform hackers

Beliefs around AI outperforming hackers have slightly increased.



Hackers who believe AI technologies increase the value of hacking

AI has already proven its value to hackers.



The AI Hacking Advantage

AI is everywhere in tech, and hacking is no exception.

But beyond the hype...

What's actually happening?

How are hackers using AI?

What does it mean for security?

Who are the AI hackers?

Nearly everyone—about 82% of hackers—is already using AI as part of their workflow. And for those who aren't using AI yet, most plan to start soon.

When we look at the breakdown, more experienced hackers are slightly more likely to have adopted AI tools, which makes sense. Hackers who have been doing this for a long time know what's tedious and what's creative, so they know which components to automate.

” HACKER AI PREDICTION

In 2026, we will see the consolidation and mass adoption of cybersecurity tools powered by LLMs. These tools already exist, but they remain expensive and experimental or require significant manual integration. Next year, plug-and-play solutions will emerge, accessible and designed for real teams, covering everything from automated application analysis to recon, low-complexity exploitation, and code review assistance.

Lower token costs will enable organizations of all sizes to adopt AI as a continuous security assistant rather than an experiment.

AI will not replace the role of the hacker or security consultant; on the contrary, it will amplify their work, increase their speed, and allow them to focus on complex findings while delegating repetitive or high-volume tasks to agents.

2026 will be the year when AI stops being an 'add-on' and becomes an essential layer within the processes of prevention, detection, and vulnerability discovery for companies, pentesters, and bug bounty hunters alike.

bronxi

How do hackers use AI?

Hackers are using AI to automate menial tasks, save time, and get unstuck. Here's how.

Speed and automation

Hackers use AI to scale their operations, such as generating reconnaissance tools, automating workflows, and creating custom scripts in seconds instead of hours. As one hacker puts it simply, **"AI automates the boring stuff to save speed and time."**

AI also accelerates learning and problem-solving. Hackers use AI to understand new frameworks, debug errors, generate payloads, and explore different methodologies. "Instead of wasting time digging through endless documentation or testing commands manually, I let AI quickly generate payloads, help me understand obscure errors, and simulate attack scenarios," one hacker explains.

Code analysis

AI excels at parsing code humans don't want to touch, like messy JavaScript, unformatted files, and massive codebases. Hackers feed these files into AI tools and get back summaries, vulnerability assessments, and analyses that would take hours to produce manually. This allows hackers to review far more code than before, finding vulnerabilities in areas that would have been skipped due to time constraints.

39 **AI creates custom scripts that allow me to avoid repetitive tasks and save more time for hacking. Also, it gives me different ways and methods of approaching a target. This way, I'm working alone, but at the same time, it's like I'm working with a team.**

anonymous

Getting unstuck

When hackers encounter unfamiliar tech, AI becomes a research assistant that can remove obstacles.. Previously, they would go down rabbit holes, wasting time, but now they can maintain momentum.

Another side effect: Hackers stay more motivated when they spend more time making progress rather than getting stuck. As one hacker explained: **"I can move quicker, which makes me want to keep going."**

39 HACKER AI PREDICTION

In 2026, hacking is going to speed up thanks to AI. We're already seeing people use LLMs to dig through data faster, spot patterns, and take on the boring parts of research.

This trend will grow, with more hackers using AI to cut down the time it takes to figure out where to look and what's worth poking at. The big shift won't be brand-new tricks but how quickly someone can go from an idea to action with AI doing the heavy lifting behind the scenes.



The art of hacking with AI

**AI just amplifies your own abilities.
What you put into it is what you get back.**

As hackers widely adopt AI tools, some are wondering: Does AI strip the art out of hacking? According to hackers, no. 95% of hackers believe hacking is an art form—regardless of whether they use AI.

That's because AI doesn't replace creativity. Instead, it handles the manual, boring tasks, such as parsing scan results or generating boilerplate code, so hackers can spend more time on the creative aspects that require their attention.



95%
of hackers think
that hacking is
an art form.

How hacking with AI helps organizations

So what does this trend mean for companies working with hackers? In short, AI changes how hackers work, which changes what they can deliver.

Faster findings

One of the biggest time sinks in security research is sorting through noise. Scan results, reconnaissance data, and potential vulnerabilities all need to be filtered, prioritized, and validated before real testing can begin. By letting AI handle this grunt work, hackers spend less time staring at logs and more time finding actual vulnerabilities. As a result, organizations get faster results, higher-quality findings, and reports that focus on what actually needs fixing.

As one hacker explained, "I use AI for automated triage. It cleans scanner noise, removes duplicates, assigns severity levels, and suggests verification steps. This makes reporting faster and clearer."

Mostly I use AI to speed up reconnaissance and triage. Instead of manually reading through pages of scan results, config dumps, and logs, I paste the noisy output into an AI prompt and ask for a concise summary: what looks unusual, which findings are likely false positives, and which items deserve hands-on testing first.

That saves hours of repetitive work, surfaces the most promising attack paths, and helps me focus my creativity on the real problems rather than busywork — plus it makes writing clear, prioritized notes and reports much easier.

More testing for the same budget

AI functions like an additional team member, handling tasks that would normally require another set of hands. This means hackers can test more attack surfaces, explore more edge cases, and delve deeper into complex vulnerabilities—all within the same engagement timeline.

For organizations, this translates to more thorough assessments without expanding headcount or budget.

”

“AI gives me different ways and methods to approach a target,” one hacker shared. “This way I’m working alone, but at the same time it’s like I’m working with a team.”

anonymous

Broader security coverage

AI can now handle tasks that were too complex or time-consuming to automate with traditional scripting.

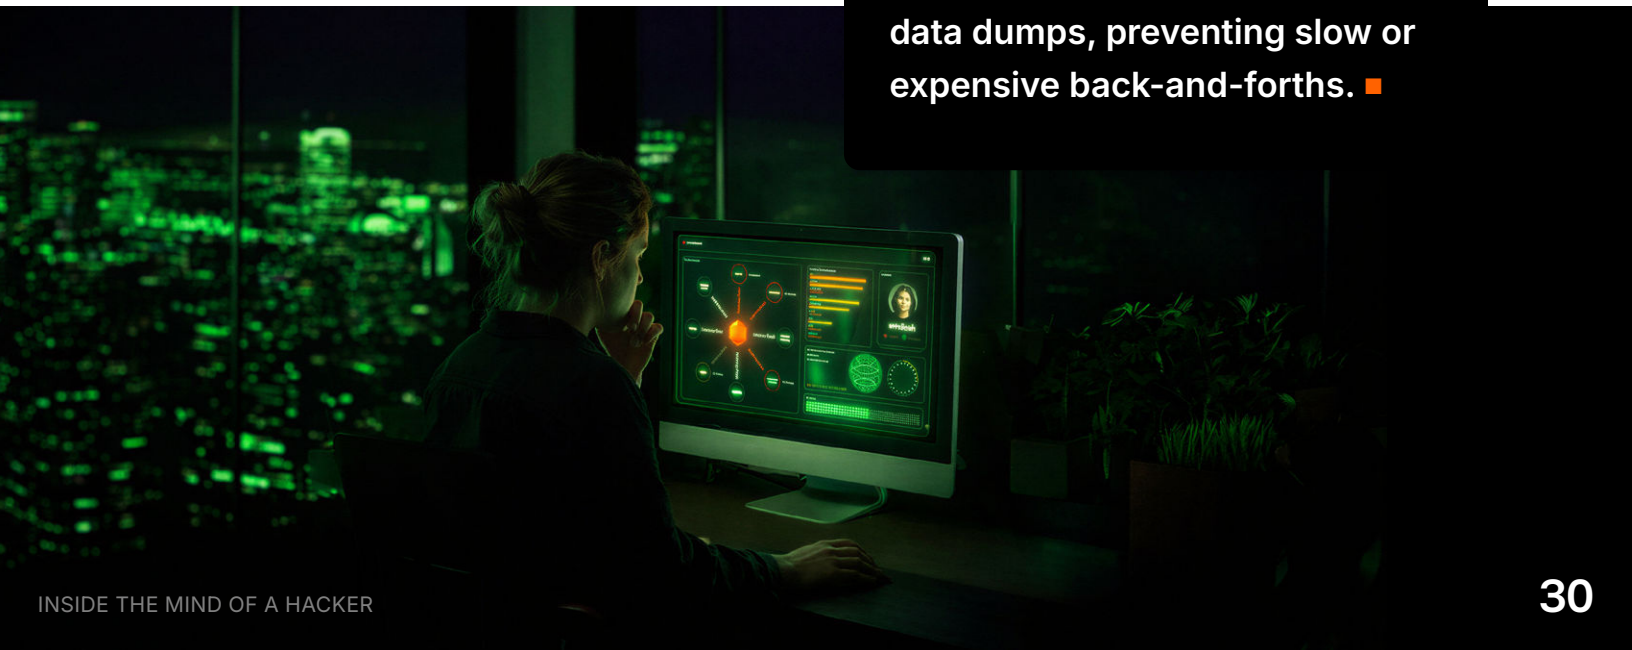
This means hackers can build custom tools tailored to specific targets, analyze obfuscated code at scale, and test edge cases that would have been too tedious to explore manually.

For security teams, this translates to more comprehensive security assessments from the same engagement. More attack surfaces tested, more edge cases explored, and vulnerabilities found in areas that would have been deprioritized due to resource constraints.

Higher quality reporting


AI helps hackers draft vulnerability descriptions, document impact, and polish their findings. As one hacker put it, “AI helps me polish vulnerability reports for better submission quality.”

As a result, companies get better reports and actionable intelligence, rather than raw data dumps, preventing slow or expensive back-and-forths. ■



Meet Aaron Guzman

CISO, Cisco Network Product Engineering



Aaron Guzman is the CISO of Cisco Network Product Engineering, a **leading organization securing enterprise and industrial networking devices from wireless access points, routers, switches to IoT cameras and sensors**. The company plays a pivotal role in securing the infrastructure that moves the world's data. Aaron is the author of IoT Penetration Testing Cookbook and has served as technical reviewer for Practical IoT Hacking and Bug Bounty Bootcamp. As OWASP's IoT Project Leader, he leads the IoT Security Testing Guide initiative.

Like so many security leaders, Aaron started as a hacker, fueled by a curiosity to take things apart, understand how they work, and make them do things they weren't designed to. The curiosity that once drove him to take apart toy cars to see what made them move now applies to hardware, firmware, supply chains, and software at enterprise scale. We sat down with Aaron to understand how hacking and security leadership go hand in hand, the challenges of securing hardware, and the impact of AI.

Leveraging the hacker mindset as a CISO

Aaron's hacking origins are rooted in hands-on problem-solving and learning to figure out how systems behave.

"Growing up without internet service at home made connectivity feel like a privilege, not a given."

"This led to me hacking access points and wireless networks, not maliciously, but because connectivity was valuable and I wanted to understand how it worked."

Aaron later applied these values, along with a strong work ethic, to several technical support roles, followed by application security and pen testing roles. Finally, he found his way into leadership, but the hacker mindset never left. “I followed my curiosity through networks, systems, communities, and eventually into the rooms where decisions get made about how organizations defend themselves. Now, instead of breaking into systems, I break down organizational assumptions about security.”

The question changed from ‘How do I get in?’ to ‘Why do we keep building things that let adversaries in?’”

Overcoming hardware security challenges in the age of AI



Hardware security is something that many CISOs are beginning to consider, but Aaron has always centered hardware as a major focus for his team. “Hardware lives at the edge of innovation, and that edge cuts both ways. We prioritize hardware security because network devices are foundational infrastructure. They’re not just assets—they’re the highways that make connectivity possible.”

“A compromised router isn’t just a singular problem; it’s a pivot point to everything flowing through it.”

This is especially relevant in the age of AI, as these hardware devices are what makes AI possible. These devices move the data that trains models and serves inference at scale, as well as the next wave of quantum computing and networking.

Aaron notes that securing hardware comes with some unique challenges:

Complex supplier ecosystems

Modern devices involve intricate manufacturing workflows and supply chains with multiple vendors, geographies, and handoffs, all of which introduce potential issues.

The debris of development

One of the most common exposure patterns we see comes from the test-and-debug interfaces left behind in production. These exist for legitimate development and manufacturing reasons, but when they ship to customers, they become part of the attack surface.

Hardware defects don't patch easily

When you find a fundamental flaw in hardware design, you're looking at redesign, remanufacturing, and redistribution as the only remediation options. The operational impact and cost are massive, as recalls, replacements, and customer disruption are often involved.

Fragile upgrade paths

The quality and reliability of upgrading low-level firmware elements determine whether products succeed. Unreliable, buggy products don't sell. For network infrastructure, the stakes are even higher because these devices run critical infrastructure that requires resilience.

Overcoming these challenges requires reimagining foundational design and quality practices for the era of AI-powered adversaries.

“The attackers of tomorrow will find vulnerabilities faster, chain exploits more creatively, and scale attacks more efficiently than we've ever seen. Our security practices need to evolve ahead of that curve. This means minimizing exposure through design—not just testing. This means treating security and quality as inseparable. This means embracing collaboration, transparency, and partnership across the ecosystem.”

“Better outcomes come from working openly with hackers, suppliers, and partners rather than treating security as a proprietary secret. The devices we build power the world's connectivity highway. That's a responsibility worth taking seriously and an incredible honor to have.”



Aaron believes that AI creates massive value as an accelerator, both offensively and defensively. He's actively finding methods to measure its effectiveness across security testing for devices as models evolve.



“New agentic capabilities are emerging that decompose security processes into automated workflows with organizational context, completing in minutes what previously took hours or days. Skills—an open specification (agentskills.io) that enables AI agents to acquire new capabilities through portable, reusable instruction sets—are being adopted by leading providers, enabling interoperability from personal workflows to enterprise security testing. This fosters a dynamic partnership between AI agents and humans,” Aaron says.

“The hackers who learn to collaborate effectively with AI will find more bugs, produce better reports, and ultimately have greater impact—whether through bounty earnings, full-time security roles, or shaping how organizations defend themselves.”

One area where Aaron envisions huge opportunities is augmenting internal pen test teams with the assistance of AI. He believes teams must embrace AI security testing, shifting the model from “humans doing all the work” to “humans as operators managing agentic workflows.” This dynamic partnership between AI and humans where AI handles scale and speed while humans provide judgment and direction is critical. “AI alone won't replace the human creativity needed to find novel attack paths, but a little nudge from AI definitely helps. I suggest teams start with small, impactful workflows and build on them. It's only going to get better from here,” Aaron declares. ■

Top Ways Hackers are Using AI

By OxMoose

+ 🌐 🔊 📄 5.2 Instant



For most hackers, the rise of LLMs, which are a subset of machine learning and thus part of the broader AI landscape, have significantly excelled our workflows.

They've changed things from "I do everything solo and manual" on late-night caffeine-fueled rampages to "I work with a paired, always-on, and enthusiastic context-aware sidekick."

For some bug hunters at least, the rise of AI isn't about replacing the human intuition required to find weird edge cases. It's about reducing clutter, fostering intuitive pattern recognition, and accelerating the parts of the methodology where machines shine: pattern expansion, contextual synthesis, and structured report drafting. In my case, I'm leveraging the power of open-source (OS) and on-demand models so that I'm not spending a dime or shipping my findings to upstream inference providers. Here is a breakdown of some of my biggest use cases for AI.

Report writing

One concrete example and likely the most popular is report writing, aka every hacker's painful woe. I built human-in-the-loop (HITL) multimodal agents that swarm on my mojo as assisted intelligence.

They capture inputs, annotate screenshots, PoC recordings, and generate clipped network requests to stitch my quick-hitting notes into a structured report designed for your eyes. The agent's role is to transcribe, tag, and suggest a first-draft report featuring an objective summary, an impact assessment, reproduction steps written for multiple audience types, and an addendum with technical artifacts. I use man-in-the-middle (MITM) approval processes, never YOLO mode, or "copypasta."

Not only does my agent make the whole process quicker, easier, and less tedious, it indexes exploits and loads notes into a vectorized database.



Think of this as a well-oiled bug hunter's vault of thoughts and a critical step that allows me to set up alerts. Specifically, my agent hands me context on wins from months ago while I'm pursuing new targets so that I can scale my exploits by demonstrating the same or varied attack strategies and techniques across web applications and enterprise infrastructure.

Sometimes, it can even help with regression testing. Yes, this is something I often monitor across your ecosystem, and you'd be stunned at how often it works.

With high precision, multimodal capture, and the ability to load months of my late-night bat signals into its RAM, the agent learns and shifts my attention to prioritize what matters most to stakeholders: the screenshot that shows the vulnerable element, the clip that demonstrates the timing, the console log that contains the error, and the final payload or bypass that successfully marked the difference between an informative report and an impactful valid bug. The agent annotates such assets and suggests which belong in the executive summary versus the technical appendix, helping deliver immediate value to the crowd. The magic isn't so much in the model inventing novel attacks but in its helping me communicate precisely and defensibly what I observed.

For non-native speakers or polyglots like my friends, LLMs can act as a linguistic bridge, offering translations that go beyond simple word-for-word conversions. They can highlight language subtleties, suggest idiomatic expressions, and even provide explanations on concurrency considerations or common pitfalls in specific programming languages. This capability allows bug bounty hunters to write more efficient and effective write-ups, ultimately leading to better-crafted, business-contextual exploits and more impactful findings.

This approach does three things for a Bugcrowd customer:

- It increases report clarity so triage teams can reproduce faster
- It standardizes linguistic style across researchers, reducing back-and-forth
- A richer paper trail of how the finding was discovered and verified with screenshots and timestamps aligned with narrative.

Let's talk augmentation using AI

There's been curiosity, sometimes insistence, that LLMs can generate payloads on demand and act as an automated exploitation assistant who's always incredibly eager to push the CVSS scale toward the ceiling.

Here's my pragmatic take ↓

Models are useful for hypothesis generation, producing "hot takes" and inspiring me toward achieving max payouts. However, handing them a live target arguably crosses a line until AI activity reaches a state of being fully authorized, auditable, controlled, and evaluated.

When it comes to my domain object model (DOM)-aware assistants, they capture the entire client-side context of most huge web applications, detect unexpected behaviors, and influence next steps with crafted specialized payloads at scale. What this means is that I can pop out some cool bugs. In other words, the agent's capabilities highlight high-signal areas and chain gadgets across probable attack surfaces with creative, critical thinking.



In the context of bug hunting, my agents augment my work by producing sophisticated checks or mental models. They offer a checklist of common misconfigurations to verify, a description of why a class of vulnerability matters, or a red team-style threat narrative you can use in a report. These outputs speed up experiments, reduce human cognitive load, and manage crucial steps in my flow that could make the difference between a P3 and P1 finding in your scope.

AI-generated code

The software development industry is locked in an arms race to produce novel features and churn out code at breakneck speeds.

This relentless pursuit has given rise to the burgeoning practice of AI-driven coding, where LLMs are harnessed to assist or even generate substantial portions of software. While this innovation holds immense potential for accelerating development cycles and democratizing coding expertise, it also introduces a new breeding ground of vulnerabilities sitting and waiting for me and others to find. The types of vulnerabilities I am identifying have changed dramatically since this particular shift in the software landscape; I'm finding all-time record counts of broken authentication, injection attacks, or remote code execution. It's practically guaranteed that I come across sensitive information outputs not limited to misconfigured security settings and/or credential leaks. Heck, take into account the recent vote on the OWASP API security project's top 10, which reflects the same results.

As ethical hackers, we must embrace this new frontier by incorporating AI-generated code into our vulnerability discovery arsenal.

By staying abreast of the latest advancements in LLMs and understanding their limitations, we refine our methodologies to identify and look for the cracks in our network proxy.

Reinforcement learning and contextual speed

In the thrilling world of bug bounties, I've found myself harnessing the power of reinforcement learning to turbocharge my agents and amp up my game.

Picture this: I'm not just any hunter anymore—I'm a maestro conducting an AI orchestra, fine-tuning each digital minion to play symphonies of vulnerability discovery.

My reinforcement learning agents, those clever critters, are like eager puppies who go after the juiciest bones. Through trial and error, they explore labyrinthine codebases nudged by my guiding principles that reward successful exploits and penalize dead ends. With each cycle, they grow more adept, refining their understanding of what makes software tick (or, sometimes, what makes it tick like a time bomb).

These AI apprentices are a new cheat code, sifting through mountains of context at breakneck speeds and unearthing complex vulnerabilities that'd give even the most seasoned hunter pause.

By constantly iterating and improving their strategies, these agents rapidly advance my own capabilities too. I'm no longer bound by the limitations of human fatigue, the types of vulnerabilities specific to my current skill set, or caffeine-fueled exhaustion. Instead, I'm propelled forward by an ever-evolving AI squadron, each one a miniature Einstein of exploit discovery. And as they learn and adapt, so do my bonds that grow stronger with every byte cracked and bug claimed.

LLMs as a teacher: Learn faster, code smarter

I love putting LLMs into teacher and pair programming tutor roles so that I can benefit from AI as a productivity tool and truly make every day a "school day."

This ease of learning is a gift from the GPU gods and something I feel a lot of people take for granted. As a nonlinear learner and somebody who dropped out of school, I wish these tools had been available to me earlier. For example, when I'm switching languages, say, rewriting a tiny recon utility from Python into Go to gain performance or compile-time portability, an LLM could make a huge difference. I don't expect the model to write the finished, production-grade scanner for me. What I want is a clear, human-readable translation that highlights language differences, concurrency considerations, and common gotchas. If I'm using an OS project in my hacking flow and desperate for a feature to springboard my skills, I'm also no longer limited to existing forks or contributions, regardless of the coding language or technology stack.

Moreover, OS projects such as Caido or Burp Suite plugins, custom vulnerability detection tools, automated bug-reporting templates, or AI-assisted reconnaissance scripts can significantly enhance a hunter's reputation and visibility within the cybersecurity community.

These contributions not only enrich collective knowledge but also provide practical benefits such as networking opportunities, mentorship prospects, and potential collaborations with other industry professionals.

The unforgiving nature of bug bounty success

In this relentless game, success isn't handed out; it must be earned through unyielding dedication to continuous learning.

The bug bounty landscape is a battleground where the most astute, experienced, and technologically advanced competitors thrive. Every second counts in this high-stakes gauntlet, and staying ahead of the curve is paramount to success and mental well-being.

Right now at least, regardless of how AI might enhance our abilities as hackers, it's crucial to remember that these tools are only enablers; they don't replace the fundamental human elements of creativity, resilience, and relentlessness in the face of failure. If we're not leveraging cutting-edge technology like AI, we risk falling behind. It's not just about keeping up, it's about setting the pace. AI can be our ally in this journey, but it's ultimately up to us as individuals to ensure that we're always at the forefront, pushing boundaries and reaping the rewards of our tireless efforts.

In essence, security is a race where the finish line keeps moving and the only way to stay ahead is to run faster than everyone else. By utilizing AI, the pack not only gets "smarter" but also quicker, running on a relentless 24/7 schedule. ■

smarter

quicker

Conclusion


For the upcoming cohort of next-generation hackers especially, AI serves as a living knowledge repository, providing instant access to vast amounts of information across diverse topics. From understanding obscure coding standards and nuanced vulnerability classes to learning about the latest exploitation techniques and industry best practices, these models help bridge the gaps in our collective understanding.

This wealth of readily available information accelerates the learning process, sets up record-breaking talent to raise the bar, and enables hunters to quickly adapt to new domains, technologies, or methodologies relevant to offensive security testing.

The symbiotic relationship between humans and AI in offensive security testing represents a powerful alliance spurring unprecedented progress in cybersecurity research and practice. By harnessing the strengths of both domains—human creativity, critical thinking, and deep domain expertise complemented by AI's computational prowess, pattern recognition, and scalability—hackers can fortify their craft, elevate individual and organizational security, and reimagine the very nature of ethical hacking.

The future of cybersecurity, powered by AI, promises a world where it is not just about defending against threats but preemptively shaping a resilient digital landscape for all to thrive within. It's time to embrace human-augmented intelligence.

Learn more at www.bugcrowd.com



INSIDE THE MIND OF A HACKER

Take what you need

A fresh zero-day

An AI agent to the dirty work

A threat model

Coffee before chaos

Curiosity (use wisely)

One clean exploit

Patience for fails

A hacker mindset

bugcrowd