



Common AI misconceptions debugged!

BY GREG JENKINS · APRIL 17, 2026 · LAST UPDATED ON APRIL 23, 2026

What you will learn

- How AI is boosting researcher productivity
- How new researchers are approaching bug bounties
- Why the quality of submissions is not declining
- How effective triage and coordination are crucial

AI and the growing ecosystem of tools built around it have now moved beyond early experimentation and into everyday use across the bug bounty community. What initially showed up as [AI-written reports](#) has evolved into something broader: changes in how researchers work, how submissions scale, and how programs experience that volume.

In the first part of this series, we explored [how AI is shifting the fundamentals of bug bounty](#), from discovery economics to operational pressure, and why the model itself needs to evolve to keep pace.

This article continues that conversation by bringing in another perspective on how these changes are playing out in practice. As with the first piece, this series brings together voices from across the company who encounter these shifts daily, spanning triage, researcher engagement, and program operations.

For this next perspective, we turn to Greg Jenkins, who leads Product and Data with a focus on turning market reality into durable product direction. Over more than two decades managing products across the tech landscape, spanning e-commerce, fraud, and telecom, Greg has built a reputation for connecting what customers need, what the business is trying to achieve, and what the platform must become to win.

At Intigriti, Greg is responsible for guiding the product against the company strategy, keeping teams anchored on the problems that matter most: the biggest challenges facing customers running security programs, and the hackers whose work makes those programs valuable. That often means leaning into forward-thinking models, not just improving today's workflows, but anticipating how the bug bounty market and researcher behaviour will evolve, and designing the systems, incentives, and data signals needed to stay ahead.

That is also why Greg's voice is essential in the "AI era" discussion. As AI changes the speed and shape of offensive security, Greg is committed to uncovering the data that reveals what is actually happening, in submissions, in triage, in researcher tactics, and in outcomes. The goal is not speculation, but clarity: insights that can shape the product and services in the months and years ahead, and help the marketplace scale with confidence.

Introduction

The term 'Vibe Coding' was coined in 2024 and went on to become a mainstream talking point not long after. Now, in 2026, we have another label doing the rounds: 'AI slop'. It's a useful phrase when it shines a light on real issues, but it's also starting to get used as a catch-all term for anything people don't like about AI, and promotes negative perceptions about AI-generated content.

At the same time, Intigriti is experiencing a surge in vulnerability submissions, and AI is a significant contributing factor. This isn't isolated. Reports suggest the impact is also being felt across other crowdsourced security companies and organisations running their own bug bounty programs. As a result, questions around quality naturally follow. One of the most common assumptions we're seeing is that bug bounty submissions are getting worse because researchers are using AI more than ever.

That's not what our data shows. Rather than a decline in quality, it points to structural shifts in vulnerability research increasing scale and velocity.

So, let's frame this as three trends and what they mean for coordinated bug bounty models going forward.

Trend 1: Researchers are becoming more productive (and AI is a meaningful driver)

At a high level, AI is doing what technology usually does when it's genuinely useful: it's compressing time.

For a researcher, that might look like:

- Faster recon and note-taking
- More rapid exploration of attack paths
- Quicker iteration on payloads and proof-of-concepts
- Easier translation between ideas, tooling, and reporting

The practical outcome is simple: in the same amount of time, a researcher can test more and submit more.

What does this look like in practice?

Historically, a new researcher starting out might achieve ~15% validity. In other words, only one or two of their first ten submissions are accepted, because they're still building skill and intuition.

That baseline dynamic hasn't changed.

What has changed is the pace. If AI helps that same person move faster, they don't necessarily become "more valid" overnight, but they may now submit five reports in the time it previously took to submit two.

The ratio does remain stable, while volume increases.

Making the assumption explicit (and why it's still a reasonable one)

To be clear: we can't "prove" from submission data alone that AI caused the increase in volume.

What we can do is be explicit about the assumptions we're making:

- We know AI tooling for vulnerability discovery matured rapidly through 2025.
- We've seen a wave of releases, workflows, and content that make offensive security more accessible.
- We can reasonably infer that this increases throughput for many researchers, especially those newer to the craft.

To strengthen the narrative (and avoid any implied leap), we can tie the submission growth curve to a timeline of notable AI-for-appsec developments or annotate the chart with those milestones. This helped create a clearer view of the events that have shaped the industry and current trends.

A quick preview of the bigger conclusion

If AI is increasing researcher throughput, then crowdsourced bug bounty doesn't become less relevant. It becomes more important. Because coordination is what turns volume into signal. This means deriving insights and value from the context being created.

And if the ecosystem is structurally moving toward higher throughput, then the industry will increasingly need to adopt technology throughout the full lifecycle, including triage, validation and prioritization, to keep the model scalable.

Trend 2: We are attracting new researchers faster and in bigger numbers (and AI likely lowers the barrier to entry)

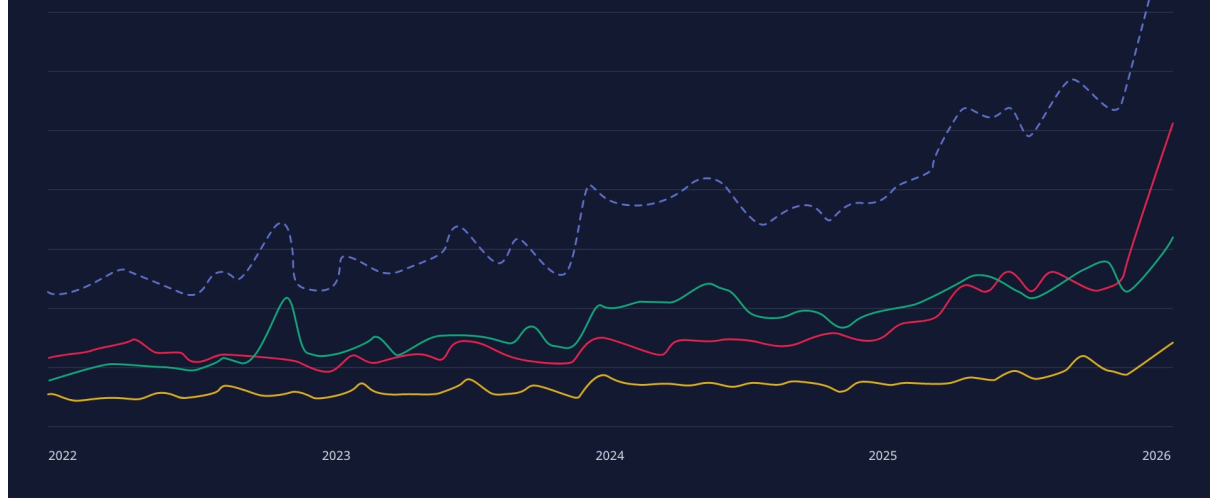
The second trend is participation.

We are seeing more first-time submitters, more sustained activity, and a faster registration rate of new entrants.

Submissions by business type

Monthly total submissions

■ Total researchers ■ Existing researchers ■ New researchers ■ Re-engaged researchers



Key: Existing Researchers: Green; New Researchers: Orange; Total: Purple

Again: we need to be clear what we know vs. what we infer.

What we know:

- First-time submitters have spiked since September 2025.
- The number of unique new researchers per month has increased.
- New researchers are producing more submissions in their first 30 days than they did a year earlier (almost double).

What we infer (and why):

- AI reduces friction in early-stage learning.
- It helps people get from “curiosity” to “attempt” faster.
- It enables researchers to iterate and submit without needing the same depth of experience at the start.

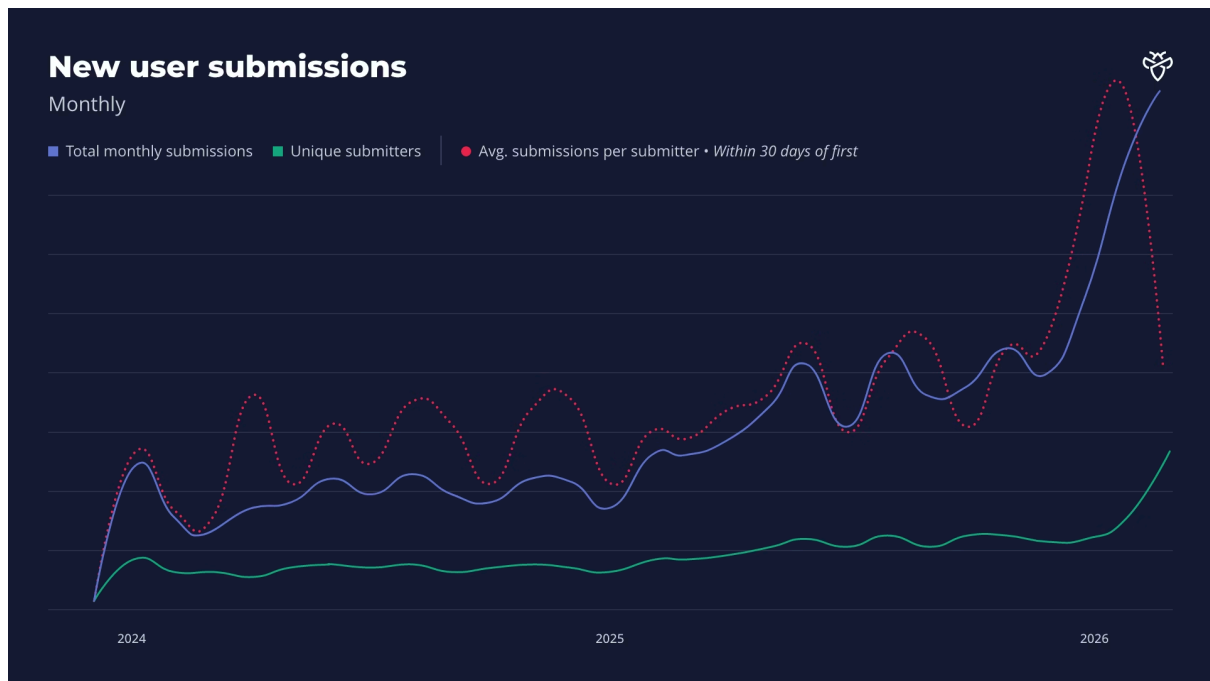
The result isn't necessarily “more skilled researchers”. It's more researchers, sooner.

Monthly new researcher submission metrics: what this implies

When you combine:

- Rising counts of new researchers
- Rising submissions per new researcher

...you get a structural shift in total throughput.



Each researcher is now **contributing almost twice** as much as they were compared with the previous year.

While this adds real value, it also creates new operational pressure, particularly in triage, as higher volumes need to be filtered, validated, and prioritized.

Trend 3: Validity ratios remain constant (AI slop isn't rising as a proportion)

AI slop is real in some corners of the internet. But the assumption that it's rising as a proportion of bug bounty submissions doesn't hold up when you look at acceptance rates.

New Intigriti findings show that while more researchers are participating and while submission volume is materially higher, the validity ratio remains broadly consistent.

In other words:

- There's more total volume
- There's more noise in absolute terms
- But noise isn't "taking over" as a share of the whole

This is consistent with what we see from top researchers too: most of the best hackers use AI to some degree and they still produce excellent findings.

What does the data show?

From 2022 to 2025, submissions grew by 328%.

Increased engagement from the researcher community is a positive signal for the bug bounty ecosystem which means more eyes reviewing programs, ultimately strengthening security coverage. The data shows we have a lot of new people coming in, and that is something to celebrate. It reflects that so many new people want to work in this space and that technology is making vulnerability research more accessible to them. As a result, we are getting more engagement from both new and existing researchers, which is fantastic for both researchers and businesses.

Since the start of 2025, submissions from new users have grown to become equal to those from more established researchers. As of February 2026, we saw a peak driven primarily by higher per-user submission rates.

And while the quality of submissions dipped in 2024, it has remained steady since the start of 2025.

Challenges (and how Intigrity is supporting the new norm)

Triage is becoming a critical part of the bug bounty model as submission volume continues to grow. It is the layer that determines whether that volume translates into meaningful outcomes or simply adds noise.

And as the ecosystem scales, the triage workflow has to scale with it. That doesn't mean "replace people". It means recognizing the nature of the work and where different approaches are strongest.

Many parts of triage lend themselves to automation:

- Some steps are deterministic and rule-based, where consistent automation can be less error-prone than human-led review.
- Some steps are judgment-heavy and context-heavy, where humans remain essential.
- And a growing number of steps sit in the middle: places where agentic workflows can bring context to bear quickly, boosting both speed and accuracy.

How AI decision support brings context to bear

A single triager, even an excellent one, can only hold so much context at once:

- Past duplicates across thousands of submissions
- Known program specifics and historical edge cases
- Standard exploit patterns and vulnerability taxonomies
- Customer expectations and policy nuance

AI decision support can widen that context window. Done well, it becomes a "second brain" for triage:

- Surfacing similar reports and likely duplicates
- Extracting key signals from long or messy submissions
- Suggesting next questions to ask the researcher
- Guiding reproduction steps based on patterns

- Helping reviewers consistently apply policy and scope

What this looks like in practice at Intigriti

Our AI Triage Assist (validation tool) can automatically scan through submissions to support triage and, therefore, customers. Triagers still review submissions but instead of starting from scratch, the AI is suggesting “this looks invalid”, and our expert human-in-the-loop makes the decision and performs the final check.

The advanced submission search allows our team to quickly identify similar or duplicate reports within submission content, helping speed up triage decisions.

We’ve already seen a significant increase in what triagers can do, just from Triage Assist. But the key point is that human experience and expertise remain central, while technology reduces the cognitive load and helps our teams spend more time on judgment and less time on repetitive scanning.

This allows us to maintain fast and consistent triage, even as submission volume grows, without compromising on quality or decision-making.

Key takeaway and looking to the future

The key takeaway from the data is that the industry is booming and AI is amplifying that momentum.

But it isn’t breaking bug bounty. It’s increasing throughput.

And that reinforces two conclusions:

- Coordinated bug bounty models are more important than ever because coordination is how you maintain signal at scale, and this is where effective triage plays a critical role.
- Technology adoption isn’t optional, it’s necessary across the full workflow, triage included, if we want to keep pace with the new volume.

To address the challenges that come with higher submissions, we have features in the pipeline focused on scoping, reporting, and education.

We’ve already started to release and will continue to roll out product updates and platform improvements designed to reduce noise in the queue, help triage teams identify high-quality reports faster, and support researchers (including elements such as AI Scope Check).

And while AI Triage Assist is already available to support today, keep an eye out for future adoption announcements in our Triage Team newsletter and on socials.



AUTHOR

Greg Jenkins

Greg Jenkins is Head of Product at Intigriti, with a diverse leadership background and two decades of experience scaling product teams across the tech space. Having worked across telecommunication, e-commerce, fraud and music, Greg brings deep customer empathy to how B2B SaaS products are built and improved. He oversees the entire product lifecycle and data, modernising delivery for the AI era without losing focus on outcomes, and is known for building high-performing product cultures.

REQUEST A DEMO

intigriti.com/demo

VISIT THE WEBSITE

intigriti.com

GET IN TOUCH

hello@intigriti.com