



How Triage Assist is raising the bar in crowdsourced security

BY STIJN BOGAERTS · MAY 21, 2026

AI is changing the volume and accelerating the pace of vulnerability submissions.

If you've been following our recent AI series, you already know that submission growth isn't a quality problem; it's a coordination problem. As Head of Triage, Lennaert Oudshoorn, outlines in his recent post, ['The AI impact: A triager's perspective'](#), the security industry is experiencing a surge in vulnerability discovery and a relative scarcity of triagers. Validating, classifying, and sorting a growing volume of reports and applying human expertise at scale have become key priorities.

That's what Triage Assist is built to support.

What is Triage Assist?

Triage Assist is a suite of AI tools embedded in the Intigriti platform, operated by our Triage Team. It's part of a broader direction on AI-leveling with the future hacker.

The framing matters: Triage Assist doesn't decide. It proposes, and humans decide. Every suggestion is reviewable, overridable, and feeds back into a structured improvement loop.

How Triage Assist supports our triage team

Under the hood, Triage Assist is powered by our proprietary agentic workflows: systems that can reason over rich platform context (program scope, submission history, and triage signals) to produce a clear recommendation with supporting evidence. The model's performance improves over time as we learn from review outcomes and edge cases, but it already delivers our highest level of accuracy when determining whether a submission is out of scope.

Triage Assist targets two triage tasks with the highest return on investment:

1. **Duplicate detection:** surfaces likely duplicates by analysing submission content and historical signals, so triagers aren't re-reading the same finding twice.
2. **Out-of-scope detection:** flags submissions that appear outside program scope, with supporting reasoning so the triager can agree, override, or escalate.

Both features are designed to shrink the time spent on first inspection so triagers can concentrate on the submissions that actually need validation and expert judgment.

What this means for security teams, triagers, and researchers

If you're running a Bug Bounty program with Intigriti, Triage Assist means faster, more consistent initial triage and better scalability when submission volumes spike. Duplicate and out-of-scope noise is caught earlier, so your security team spends more time on findings that matter.

For our triagers, it means more time on complex, high-signal investigations. For researchers, it means faster feedback loops and fewer reports lost in a queue.

Why "human-in-the-loop" isn't a buzzword here

As we scale, we are constantly evaluating where human and artificial intelligence are deployed to best effect. Our agentic workflows are built on a decade of human decisions, not just submissions, that models can bring to bear far more consistently than humans. Equally, we see clearly the limitations of current models and the importance of human intervention, particularly in dialogue with researchers and customers, and to mitigate the risk of closing submissions in error.

Security workflows require more than accuracy. Triagers need to understand why a suggestion is being made. Teams need to be able to intervene on edge cases and document their reasoning. And the system needs structured feedback, so it improves.

What makes shipping Triage Assist in a production security platform harder than it looks

Shipping AI in a production security platform is not the same as adding a prompt to a workflow. To do it safely and reliably, we've built:

- Short-lived, scoped authorization boundaries, so the agent only ever touches what it needs to.
- Observability across model calls, outputs, and sensitive data handling.
- Retrieval tooling so the model references real program context, not made-up details.
- A review-first workflow designed to support expert judgment, not create babysitting overhead.

Triage Assist is also the foundation for a reusable agent framework that will expand across the vulnerability lifecycle.

What's coming next

Triage Assist is the first step of many. As AI continues to reshape the volume and complexity of vulnerability submissions, triage is just one part of the lifecycle that stands to benefit from thoughtful AI assistance.

And we are just getting started!

We're focused on the full arc, by providing feedback to researchers before their submission is created to the moment a fix is shipped. Better validation, faster reproduction, more effective mediation, and stronger remediation support are where we're headed, always with human judgment at the centre.

The bottom line

AI is going to keep changing the volume and shape of what lands in your triage queue. Our position hasn't changed: the right response is augmentation that preserves expert judgment.

Triage Assist is our first launch in that direction. It's available now through our triage team.

For more information on our AI features and systems, view our AI [model card](#).

We've got plenty more coming as we explore how we're building Intigriti for the next generation of hackers. So, stay tuned, and [contact us](#) if you would like to know more.



AUTHOR

Stijn Bogaerts

Stijn Bogaerts is Product Manager at Intigriti, the leading Bug Bounty and ethical hacking platform that connects organizations with a global community of security researchers. In his role, he aims to provide security researchers with the best possible experience on the Intigriti platform.

REQUEST A DEMO

intigriti.com/demo

VISIT THE WEBSITE

intigriti.com

GET IN TOUCH

hello@intigriti.com