How AI is leveraged to enhance the Intigriti platform

BY ELEANOR BARLOW · DECEMBER 1, 2025 · LAST UPDATED ON DECEMBER 24, 2025

What you will learn

- **How AI improves the Intigriti platform:** Learn how AI is used to speed up vulnerability triage, detect duplicates, and improve report quality.
- **How AI is implemented responsibly:** Understand how Intigriti combines AI models with human oversight to ensure accuracy and security.
- **What this means for users:** See how AI benefits both security researchers and customers.

What is Intigriti's stance on AI?

At Intigriti, we believe AI is a powerful ally to, not a replacement of, our community of security researchers. We will use AI to empower our researchers to hunt for bugs smarter, faster, and more efficiently, while recognizing the value of human creativity and ingenuity that machines cannot replicate. By creating AI-powered tools informed by researcher and customer insights, and built on a foundation of trust and consent, we enable researchers to focus on what matters most: uncovering critical vulnerabilities faster and securing the digital world.

Our vision is to build an integrated ecosystem that guides ethical hackers toward the most impactful targets, accelerates their research, and ensures their expertise remains at the heart of every discovery.

We are not just adding AI to the platform; we are creating better opportunities for hackers and delivering better security outcomes for customers.

How does Intigriti implement AI?

Intigriti has implemented AI in many ways to help enable customers and researchers across the company. In this article, we will focus on the implementation examples of AI on three business challenges relating to vulnerability submissions.

- Our 'Dupe detection' function flags any likely duplicate submissions to speed up triage. The model achieves 95% accuracy and is retrained every six months.
- Our 'Similarity detection' function surfaces submissions that are similar to current ones and focuses on items that are resolved or negatively rejected to enable recurrent retesting insights.
- Our 'Validity detection' function predicts if a submission is valid, acting as a filter. It relies on submission, program, and researcher metadata rather than raw text, including images, sections, lines, steps, domains, etc.

These self-hosted models, all running internally, are models that we train ourselves.

How does Intigriti monitor and improve performance?

It is important to note that Intigriti:

- Monitors and mitigates bias through testing, human review, and performance metrics.
- In-house models are retrained every 6 months with updated labelled data.
- Intigriti continuously tracks performance and alerts on anomalies.
- Major updates to prompts, embeddings, or in-house models are documented for traceability.

Most importantly, AI features act as decision aids. Final responsibility remains and will always remain with human reviewers.

View the AI Model Card for more information on features, upgrades, maintenance, and ethical considerations.

What does the AI Architecture look like?

This AI data flow benefits customers by enhancing their Intigriti services. Researchers also benefit as processes leverage AI to provide a better surface and platform for communication.

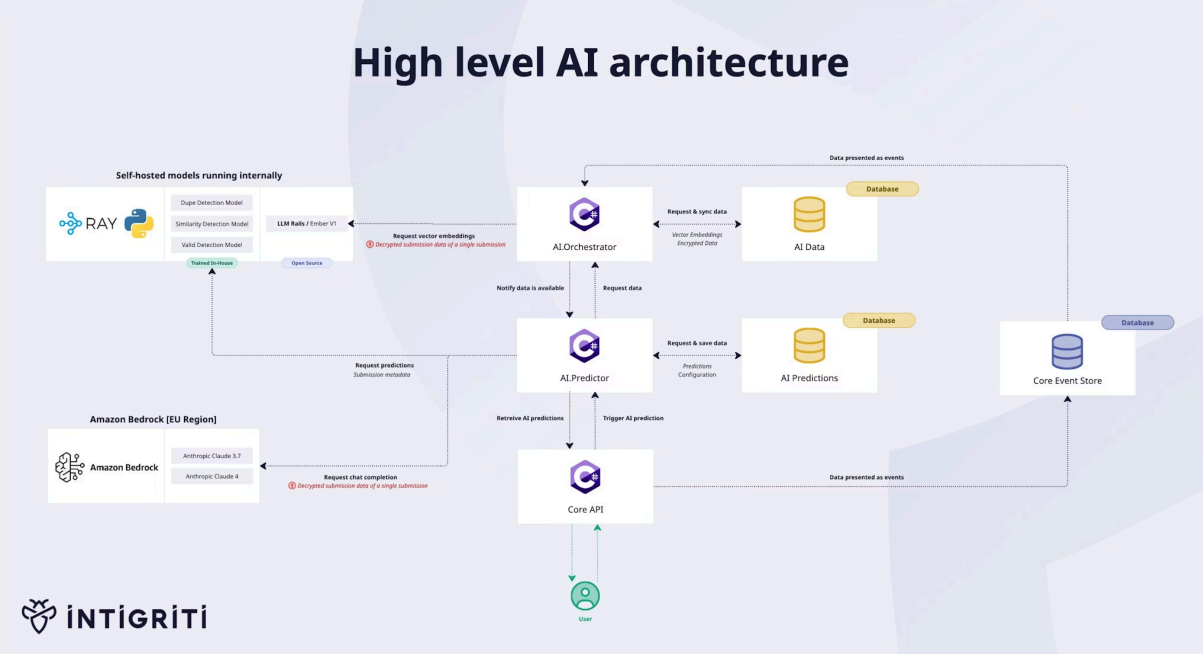

AI data flows within the Intigriti platform.

All our infrastructure is within AWS. We use a .NET based backend to orchestrate all the data flows within the platform. Located in the centre of the diagram is the C sharp programming language. C sharp encompasses different programming disciplines, including static typing, strong typing, declarative, functional, generic, object-oriented, and component-oriented.

If you start at the base, with the green ‘User’ icon, the entry point for the user goes directly into the Core API. This is where every kind of data transaction, including submissions from the researcher, takes place. Anything AI-related also goes through the Core API.

Data that is saved always goes towards the Core Event Store, and all the information is decoupled there. This means that it divides data into independent, modular components, instead of a single tightly coupled system. In essence, Intigriti tracks every change within the platform. This enables us to know what the state of data in the platform was during a specific time. When training AI models, this is very valuable as it allows us to learn from past events in past years, months, and days.

How is data handled?

The most important thing here is that data is handled asynchronously, meaning it isn't saved at the exact moment it's created. Instead, it's saved as a response to some event or action that occurred within the platform.
The AI Orchestrator is used to detect a submission-created event and will, for example, call for ‘Request vector embeddings’ to save that data in another format that we need for AI matching, and then notify the AI Predictor of the AI elements that are needed. These then filter back to the Core API and to the user.

What is the process for training data?

The training happens in Jupiter notebooks. That’s the Python element that you can see on the diagram (the yellow and blue logo).

Ray Serve, the logo next to the Python logo, provides scalable model serving for Sentence Transformers and XGBoost over gRPC, which is a Google development for open-source, high-performance remote procedure calls.

Models are usually binary classifiers, and we use XGBoost behind the scenes. XGBoost is an open-source, scalable ML library known for its speed. The XGBoost framework is retrained every six months on resolved submissions, during which the accuracy and metrics are re-evaluated.

Data from the submission, such as the following list, can be taken and fed into the model.

- Title length
- Endpoint similarity
- Number of headings shown
- Which researcher it was
- How many access submissions did the researcher have

And much more

From this, it starts training to figure out which features have the most impact, and it classifies them. The result of this model is a prediction of any kind to aid triage in swift decision-making, which makes both customers and researchers happier.

But the vector is also responsible for making the calls to our large language model, where we use Amazon Network, which provides the Anthropic Claude 4.5 (always the latest version), through AWS. We go directly via AWS for access to the large language models, which guarantees that the data that we send is not used to train by third parties.

Control levels: What models are used?

We build a scalable setup. We must work with models that are self-trained and self-hosting, which means having the ability to fine-tune a model to our own data and then running it on our infrastructure.
We use Amazon Bedrock, an AWS service, for accessing large language models. All within the AWS ecosystem.

Compliance: Where is everything hosted?

The Intigriti platform is hosted in the EU. Some third parties are located in the US, but everything is hosted in the EU-West-1 region for Amazon Web Services (AWS), to ensure compliance and minimize latency.

For Intigriti datasets, we filter out all the companies that have data AI features disabled, to make sure that we do not train on data that customers say they don't want to be trained on.

The demand for regional data hosting has increased in recent years. With anything from healthcare providers to large financial organizations and governmental processes, hosting data within the EU borders ensures alignment with regional data protection laws. This also minimizes risks associated with any transfers across borders and builds a rapport between providers, stakeholders, and organizations.

Security: What protection is in place?

Submission data is encrypted at rest, in transit, and at the application level. This means that data is protected while it is being stored on servers or devices. And it means that data is protected while moving across networks. And it is protected within the application itself. Combined, the data is protected in all its conditions.

This means that even an engineer during an incident connected to the database cannot read what's in the submission. And to do certain actions, we need to decrypt the content and have the decrypted content sent to large language models or to a self-hosted model to get the information out.

What's next?

At Intigriti, we are always looking into methods to develop and enhance our offerings and to invest in features and tools to support our researchers and customers.

> **"We believe that AI can be leveraged to support customers after vulnerability detection has been conducted by our researchers. We have some exciting solutions in the pipeline that will support in mitigating, remediating, and preventing recurring vulnerabilities."**
>
> **Arne Schoonvliet, Intigriti**

> **"Intigriti won the 2025 Cybersecurity Innovation Award for elevating bug bounty programs into a continuous, intelligence-driven security platform built around the strengths of our researcher community. We've introduced new AI-powered technology that helps ethical hackers showcase their specialised skills and get matched more effectively with the right customer challenges, improving quality, speed, and engagement on both sides. We see AI as a powerful ally to both the community and our customers, enabling discoverability to be smarter, faster, and more efficient, while recognising that, at the core, it is unique human creativity and ingenuity. The AI-powered tools we're delivering are informed by researcher and customer insights and built on a foundation of trust and consent. I'm genuinely excited by some of the technologies we're soon to be releasing."**
>
> **Gregory Jenkins, Head of Product, Intigriti**

The future of crowdsourced security will be built by those who understand that technology and people are stronger together. At Intigriti, we are committed to creating that future, one where hackers are empowered, never replaced. AI is, and will be, a powerful ally, used to empower our researchers to hunt for bugs smarter, faster, and more efficiently. By creating AI-powered tools informed by researcher and customer insights, we enable researchers to focus on what matters most: uncovering critical vulnerabilities faster and securing the digital world.

Visit this page for more information on Intigriti's AI process.

If you have a question about any of the information provided in this article, message our team here.

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.