CASE STUDY

Autonoma Deploys IoT Service Agent in Weeks Instead of Months with Vijil on DigitalOcean

PRESENTED BY

Customer Profile

Company: Autonoma

Industry: IoT for Industrial Machinery

Stakeholders: Senior technical and executive leadership

Agent: AI agent for troubleshooting technical issues

Environment: Industrial machinery service and maintenance

Constraints: Speed of development, reliability, security, scalability

The Challenge:
Rapid deployment of a reliable AI agent

Autonoma, a leading IoT platform for industrial machinery manufacturers, faced the critical challenge of addressing equipment downtime, which can be costly and negatively impacts their clients. To enhance technical support, Autonoma aimed to develop an AI agent that provides troubleshooting instructions for complex technical issues. However, the process of building and verifying this agent was anticipated to take several months.

Before partnering with Vijil, Autonoma encountered several challenges:

Technical limitations:

Early prototypes had issues with accuracy, including generating hallucinated responses and failing to switch between languages (German and English) as per user request.

Security risks:

The initial prototype was susceptible to misuse and attacks.

Protracted development timeline:

Developing custom testing, evaluation and guardrails would entail extended development cycles.

In addition, the team had decided that the agent would run on DigitalOcean's GenAI platform because of ease of use and low-touch integration easily into their existing infrastructure.

The Objective:
Uniform trustworthiness across all agents

Autonoma's mandate was clear:

Develop a reliable, secure, and efficient AI agent capable of enhancing customer service in a fraction of the usual time, ensuring the agent could accurately troubleshoot issues while adhering to strict security standards.

The Solution:
Leveraging Vijil for rapid development

Autonoma partnered with Vijil to accelerate the agent's development and enhance its reliability. The collaboration focused on several key improvements:

1. Model optimization

- The team identified that their initial model of choice generated a high rate of hallucinations when deployed with a Retrieval Augmented Generation (RAG) function connected to the firm's knowledge base - and was vulnerable to prompt injection attacks.

- Vijil security and safety testing pointed to a model that performed less well in the inclusion of images, but that would significantly lower the risk from simple agent misuse via prompt injection as well as jailbreaks.

2. Accuracy enhancement

- Vijil reorganized Autonoma's online help content, ensuring it was accessible in both English and German, thus improving retrieval accuracy for RAG functions connected to the new model - and reducing the rate of hallucinations.
- The English and German content was separated into different knowledge bases to improve the ability of the agent to cite content in the user's language.

3. System prompt refinement

- The system prompt was expanded to include strict instructions, improving agent responses and preventing misuse.
- Explicit guidelines allowed the agent to alternate accurately between languages.

4. Guardrails implementation

- Vijil integrated safety policies and misuse prevention instructions, enhancing the trustworthiness of the agent, preventing the agent from being exploited or performing tasks outside its intended scope.
- Users can now enforce these policies through guardrails deployed on the Digital Ocean platform.

The Results:
Operational efficiency and enhanced customer value

Working with <u>Vijil</u>, an expert AI agent builder on the DigitalOcean <u>Gradient Platform</u>, the Autonoma team deployed a secure, reliable agent in just one week, saving a lean Autonoma team months of development time and costs. Using the AI agent, Autonoma has reduced their customer's time to troubleshoot each incident and is poised to scale their operations faster to more customers.

The collaboration led to significant changes in the development and deployment process, achieving remarkable outcomes:

Rapid deployment

The AI agent was quickly developed and deployed in just one week, a dramatic reduction from the previous timeline.

Improved accuracy

The agent now generates reliable content with adequate references and grounding in original documentation, along with successful handling between languages.

Enhanced security

New measures made the agent resilient to common prompt injection attacks, increasing safety for end-users and reducing the risks of deployment in production.

In their words

"Working with DigitalOcean and Vijil has made it easy to dive right in and start building, without the need for a ton of resources. It also helped us to quickly deliver real value to our clients."

About Vijil

Vijil is the trust infrastructure that enterprises need to develop and deploy AI agents with reliability, security, and safety. Vijil compresses the time and effort to deploy trusted agents by 4x.