Enterprise-Grade Security, Privacy, and Accuracy

SAFEBASE

SafeBase AI

At SafeBase, we recognize the critical importance of data security, privacy, and transparency.

Our **comprehensive AI** offers enterprise-grade security and privacy, coupled with unmatched accuracy.

Quick overview of our AI pipeline

- We use a RAG-based LLM pipeline with well-known foundational models — your data is only sent to the LLM in-context for answering a question.

  In-Context: The LLM only receives your data in a prompt — it does not store or learn from it

- Your data is stored in a segregated database (per customer) in our VPC (e.g., vector embeddings)

- We don’t use your data for training purposes; our foundational model providers do not train their models on your data as well

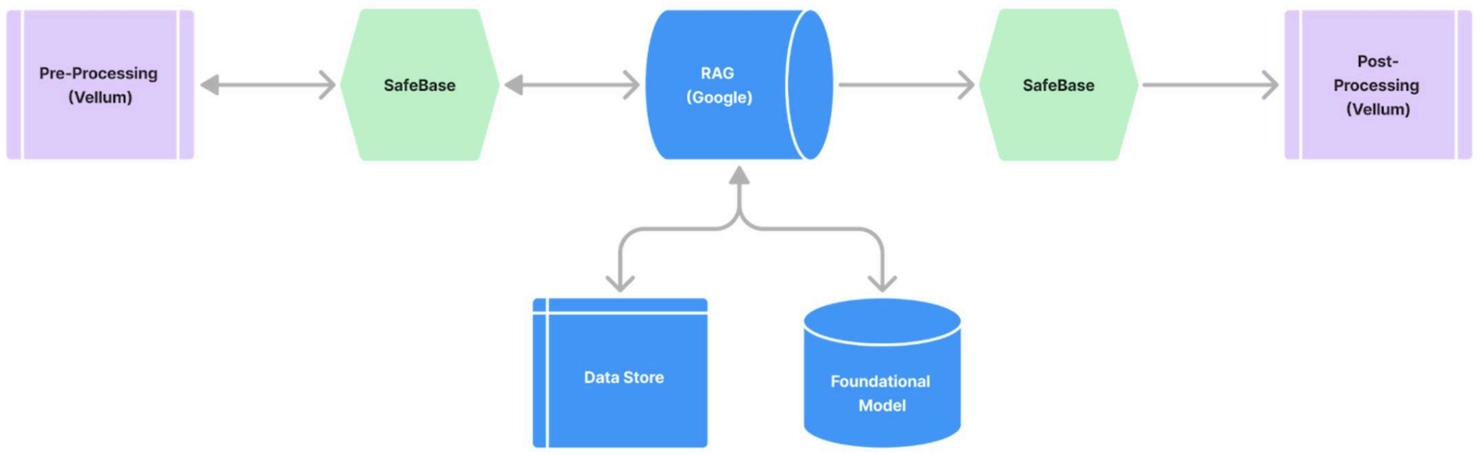

Best-in-class AI pipeline infrastructure to **optimize** customer security, efficiency, and accuracy.

In 2025, we're upgrading SafeBase AI to bring you even more benefits.

Additional functionality

Benefit from new multi-language support, index external websites to answer questionnaires, plus auto-fill questions then re-run AI answers as needed. You’ll also have new verbosity customization controls. Choose whether AI offers single word responses, single word plus additional comments, or more detailed answers.

Continuous improvements to the AI Pipeline

Our in-house team is dedicated to updating our custom AI daily — incorporating customer feedback to quicken response and integration of changes. Plus, new accuracy metrics and a robust evaluation pipeline will also drive continuous optimization.

Faster Processing

Experience even faster question response time with our upgraded AI, and ask single questions of SafeBase AI as well as full questionnaire runs.

Here's how we handle your data:

Data Storage and Processing

- Your data is stored in a DB in our VPC
- Our models can’t query your data directly, it is only passed in-context

Models and Training

- We do not use customer data for training models
- We only use approved Enterprise-grade foundational models from providers such as OpenAI, Anthropic, and Google

Currently, the main model powering our AI is Gemini *(Google)*. Our chosen model may be subject to change based on product quality & our internal assessment. In any case, we will never use a model that uses data for training

Cloud Infrastructure

- Our cloud infrastructure is powered by Google *(Vertex Agent Builder)*
- We utilize Vellum, a 3rd-party SaaS tool, to build advanced LLM workflows. Vellum only uses snippets of information to run LLM chains, it does not use your raw data directly. We do not store any data in Vellum.

Opting In and Out of AI

- AI is only used when requesting generative answers to a question, e.g., when running a questionnaire with our AI-powered Questionnaire Assistance product.
- Access to this feature is gated with RBAC to control which users may utilize it.
- SafeBase may explore other uses for AI in the future. Customers will always have the ability to opt-in/-out of AI features.

Compliance and Security

SafeBase uses a layered security approach to protect the application and customer data. Details can be found at https://trust.safebase.io/.

AI Subprocessors

- Google - data storage & foundational models
- Vellum - LLM pipeline
- The most up to date list will always be available at https://trust.safebase.io/subprocessors

FAQs

What actions are required of me and my team?

Nothing! We'll handle the upgrade in the background while you continue with business as usual. You're also welcome to inform your team about the AI enhancements.

For any inquiries or to join the beta for this upgrade:

Contact your CSM

or Macy Mody | *VP Customer Experience* | macy@safebase.io

We are dedicated to handling your data with the highest standards of security and privacy.

safebase.io