AI in Saga

Last updated: 2026-10-26.

Saga uses AI to help you organize your own data. The AI operates within Saga: proposed actions are checked by Saga and executed via its own features.

1. What does Saga use AI for?

With AI, you can give instructions regarding contacts, appointments, tasks, notes, relationships, objects, and locations. Saga can look up data, interpret instructions, and make suggestions to create or modify data. Saga can also recognize information from images and documents and process speech.

Saga is intended as a tool for your personal organization. Answers about your data depend on what is available in the vaults you have access to.

2. What data does the AI use?

Saga assembles context from your instruction, any conversation context, attachments, and data from the vaults for which you have permissions. In doing so, Saga can search across multiple accessible vaults: you do not need to know in advance where something is located.

This context may contain personal data, such as names, contact details, or appointments. The data is not automatically anonymized. The AI does not have unlimited independent access to the databases: Saga determines which context is provided and checks permissions and proposed actions before they are executed.

In case of ambiguity, Saga may ask for clarification or confirmation. These checks limit errors but do not completely rule them out.

3. Responses and execution

The model may suggest response texts with fill-in fields for an instruction. Saga fills these in based on the actual execution result. A second AI call is not made by default to embellish a response; Saga uses response templates and its own messages for that purpose.

A success message corresponds to a successful execution. If an action fails or the result cannot be confirmed, Saga uses an appropriate error or uncertainty message. The AI cannot grant rights simply by stating that something has been executed.

4. Images, documents, and speech

- **OCR (Optical Character Recognition)** means recognizing text in an image or scan. Saga combines this with AI interpretation, for example to extract a name and phone number from a business card or a date from an invitation. This processing is done with Qwen via Scaleway.
- **STT (Speech-to-Text)** means converting speech to text. When you speak an instruction, Deepgram processes the audio so that Saga can work with the recognized text.
- **TTS (Text-to-Speech)** means converting text to speech. ElevenLabs receives the text of a response when you have it read aloud.

For a PDF, Saga tries to read the text directly from the file. Where image recognition is needed, pages can be processed as images. For these functions, images, document content, audio, or text may be passed to the respective supplier. Only share the content you want to use for the function.

5. Suppliers and retention

Scaleway/Qwen: AI processing takes place via Scaleway in Paris. According to Scaleway, the content is not used for training the base models and the model maker has no access to it. There are temporary inference caches. In exceptional cases of error or abuse investigation, full requests may be stored for a maximum of two weeks. Aggregated, anonymized usage data may be stored for a maximum of six months.

Deepgram: Saga uses the EU speech route and sends the setting that excludes requests from model improvement. According to Deepgram's documentation audio and transcriptions are stored with this setting only as long as necessary to process the request. Technical request and usage metadata are not covered by this. The transcription can then become part of your conversation or stored data in Saga.

ElevenLabs: for the Saga account, use for model improvement has been disabled. This is not the same as no storage. ElevenLabs may store the provided text and generated audio, among other things for history, error investigation, and security. Saga makes no promise of zero storage or exclusively European processing for this service. See ElevenLabs' explanation on storage.

The supplier overview describes which services Saga uses. The privacy policy describes the processing of personal data by Saga. The processing of an attachment by AI does not mean that the original attachment is automatically removed from Saga.

6. Verify important outcomes

AI may misread text, misunderstand speech, confuse persons, or misinterpret dates. Therefore, check important data and changes, such as contact details, persons involved, appointments, times, and locations. Saga does not guarantee that AI output is always correct or complete.

7. Feedback and bug reports

You can report errors via feedback or a bug report. Depending on your choice, relevant conversation or technical context may be included. This is a separate processing activity in addition to the content required to execute your AI task.

See Feedback and bug reports for the explanation regarding included data, access, and retention.

8. Usage and costs

Saga records technical usage data, such as provider, model, token counts, duration, status, and estimated costs. Tokens are the text units that a language model processes. This cost record does not contain full prompts, transcripts, or answers. This is separate from the conversation content required for Saga to function and from context you deliberately include with a bug report.