FAQ: AI models and keys

From the SourceLace manual: Which AI sees what, using your organization's own AI key, and whether anyone trains on your data.

SourceLace version 0.1.0 · Made November 22, 2026
The latest version is at https://sourcelace.ai/docs

AI models and keys

Which AI sees what, using your organization's own AI key, and whether anyone trains on your data.

Does SourceLace send our data to an AI model?

From your own AI app: no. SourceLace returns results to your app, which uses its own AI provider under your agreement with it. From the assistant in the SourceLace app and from AI agents: yes, the conversation (or the agent's instructions) and the results of the tools it used go to the AI provider your admin chose. Agents on a SourceLace Runner use your own key from inside your network instead. Fields your admin hid or masked reach the AI only as placeholders, or not at all.

Details: AI and your data.

Does the AI see our sensitive data?

Your AI works without seeing your sensitive data. The right people still see the real values. Your admin chooses which fields are hidden from the AI and which are masked, and the AI gets only a placeholder or nothing. People your admin allows still see the real value in the SourceLace app. Everyone else with access to the record sees it as their own permissions in the system allow. The one exception is a photo, if your admin has turned photos on: the AI reads the picture, and data protection rules cannot check what is in it.

Details: Data protection.

Which AI does the assistant in the SourceLace app use?

By default, an AI model included in your plan. Your admin can instead use your organization's own key from Anthropic, OpenAI, Google Gemini, Mistral, Azure OpenAI or Amazon Bedrock, and can choose which model the assistant uses from those available. SourceLace itself works with any AI app your people prefer.

Details: AI settings.

Can each person choose their own AI model?

Yes, when your organization offers more than one. **Model**, under **Options** below the message box in any chat, new or old, lets a person choose the model for their own chats, for example a smaller, faster one for quick questions or a larger one for hard ones. It changes only their chats: Search everything and AI agents use the organization's model. A larger model uses more of the organization's credits per question, and on the included AI a plan may offer only some models. Your admin chooses the organization's model under **Settings**, the **AI model and use** tab.

Details: Choosing the AI model for your chats.

Does SourceLace summarize our chats, and can we stop it?

Only when a chat is nearly full, or when a person presses **Compact all** or **New chat from summary** in the chat's Context panel. A summary is one small call to your organization's AI model (counted in credits like a question, or billed by your provider if you use your own key). It is written from the chat as the AI already saw it, so values your data protection rules mask stay masked, and it is kept inside the chat, encrypted with your organization's key, and deleted with the chat. The saved chat keeps every message, and a person can undo a summary. An admin can switch automatic summaries off for the whole organization under **Chat defaults**, **Full chats**; people can still press Compact themselves.

Details: Full chats: compacting automatically.

Can we use our own AI key?

Yes, for the assistant in the SourceLace app and AI agents: open **AI model**, and choose your provider and paste your own key (Anthropic, OpenAI, Google Gemini, Mistral, Azure OpenAI or Amazon Bedrock). Your provider then bills you directly and questions use no SourceLace credits. Your own AI apps use their own provider and are not affected.

Details: AI settings.

Can the AI run in our own Azure or Amazon cloud account?

Yes, for the assistant in the SourceLace app and AI agents, if you have your own Azure OpenAI resource or your own Amazon Bedrock account: on **AI model**, choose **Your own Azure OpenAI resource** or **Your own Amazon Bedrock account**. The conversation then goes to that service in your own Microsoft Azure subscription or Amazon Web Services account and region, under your agreement with Microsoft or Amazon, and they bill you directly. Your admin needs to create a key (Azure) or an access key (Amazon) and note the endpoint or region and the model: the steps are in the admin guide. SourceLace does not provide these services itself.

Details: Azure OpenAI, Amazon Bedrock.

How is our AI key protected?

It is encrypted with your organization's key, never shown again, never returned by any interface and never logged. On a SourceLace Runner, it stays on the runner inside your network and is never sent to SourceLace.

Details: Keys: who brings each one, and where it is kept.

Does SourceLace train AI models on our data?

No. SourceLace does not train AI models on your data, and sends it to no AI provider other than the one your admin chose. How long a provider keeps requests is set by that provider's terms; with your own key, your own agreement with the provider applies.

Details: Do the providers keep or train on our data?.

Does using our own AI app cost SourceLace AI credits?

No. Your own AI apps do their own thinking and use no SourceLace credits. Credits count AI use in the SourceLace app and by AI agents.

Details: AI use: credits and questions.

What does SourceLace's own AI see when it helps with setup?

Descriptions of your setup, never your records. When drafting an agent, the AI sees what you wrote and the names and kinds of your systems. The setup assistant sees the names and types of objects and fields, short lists of a field's allowed values, and counts of values that look sensitive, never the values themselves.

Details: Setup assistant and Describe it.