IdeaLine · AI Usage Policy

Draft, pending legal review. This document is a draft while our counsel completes their review. It may change before IdeaLine opens to the public, and nothing here is final until we publish a reviewed version with an effective date. This is not legal advice. Questions? Write to hello@idealine.io.

AI Usage Policy

Last updated [EFFECTIVE DATE]

1. What this covers

This policy explains how AI works in IdeaLine, what we do to keep it safe and private, and what you are responsible for. It covers two things: **AI assist** (optional features that help draft, summarize and search your ideas) and **Sparky**, our help assistant (an AI agent, not a person, that answers questions about using IdeaLine). It is part of the Terms and Conditions, and the Data Policy has the data detail.

2. How it works

- **Off until you turn it on.** AI assist is off for every workspace. An owner or admin turns it on, and one switch in Settings turns it off for everyone. When it is off, nothing is sent to any AI provider.
- **Your request, one at a time.** When someone uses an AI feature, our server sends only the text that feature needs, with names replaced and contact details removed where we can. Your browser never talks to the model and never holds a model key.
- **The provider.** Requests go from our server through **Vercel AI Gateway** to **Anthropic's Claude** model. Both are subprocessors, listed in the Data Policy and used only for AI features in workspaces that turn them on.
- **No training.** We do not use your content to train AI models, and we use only providers whose agreements prohibit training on it. Vercel states it deletes prompts and responses once a request completes. We do not keep the text of prompts or answers: we record only who used a feature, when, which feature, and its size and cost. Anthropic's own retention is described in its API data retention documentation, and content its safety systems flag may be kept longer for safety purposes.
- **Not analyzed or profiled.** AI requests are not used for analytics, advertising or profiling people.

3. Limits that keep it safe

Each plan includes a monthly allowance of AI use, set by IdeaLine and shown to you only as credits left. Every request is also limited in size (for example, up to 9,000 characters of

input), in how many can be made a minute, an hour and a day per person and per workspace, and in how long it may take. Sparky accepts messages of up to 600 characters and answers only from our published help articles. We may change limits at any time and may pause AI features when a limit or an outside provider fails. AI features are excluded from every uptime commitment.

4. What you must do

- **Do not enter highly sensitive data.** That means health information, government identification numbers, financial account or card numbers, passwords or secrets, and similar. AI features are not designed for them.
- **Check the output.** AI can be wrong, make things up, or miss things. Treat answers as drafts. Review them before you rely on them or send them to anyone.
- **No high-stakes decisions.** Do not use AI output as the sole basis for decisions with legal or similarly significant effects on a person (hiring, credit, housing, health, legal advice, safety), and do not use it for emergencies.
- **Be open with others.** If you share AI-generated content with clients or colleagues in a way where it matters that it came from AI, say so, and follow the rules that apply to you.
- **Follow the rules.** Do not use AI features to break the Acceptable Use Policy or the law, or the model provider's usage policy.

5. Disclosing sensitive data, and responsibility

You decide what text you give AI features. If you enter sensitive data despite this policy, you do so at your own risk, and as between you and IdeaLine you are responsible for it, to the extent the law allows. AI output is provided "as is". We do not promise it is accurate, complete, lawful or fit for any purpose, and we are not liable for decisions you make or actions you take because of it. The disclaimers and limits of liability in the Terms apply.

6. What we do about safety and privacy

- Features are scoped narrowly: each does one job with a fixed instruction, a fixed output size and its own limits.
- Model answers are shown as text, never run as code or as web page markup.
- Sparky can link only to articles in our own help center.
- We watch cost, errors and abuse without reading your content. An owner or admin can switch AI off for the workspace at any time, and we can pause AI features for everyone if a limit or an outside provider fails.
- Questions and incidents go through our normal security and incident process.

7. AI laws

We design IdeaLine's AI to be a limited-risk, general-purpose productivity assistant. It is not built for, and you may not use it for, uses that laws treat as prohibited or high-risk, such as social scoring, emotion recognition at work, biometric identification, or automated decisions about employment, credit, education or essential services.

- **EU AI Act.** We tell people when they are talking to an AI (Sparky is labelled as an AI agent), we label AI output in the product as an "AI draft" for you to review, and we provide this information. We train our team on how the AI features work and what they should not be used for. Our role is that of a deployer of a third-party model; the model provider is responsible for its model.
- **US state laws.** We follow AI disclosure and transparency laws that apply to us, including those of Utah, California, Colorado and Texas, as they take effect. We do not use AI to make consequential decisions about people.

[Counsel: confirm which obligations apply to IdeaLine as a provider and as a deployer, and update this section.]

8. AI providers (subprocessors)

Provider	Role	What it receives
Anthropic, PBC	AI model provider (Claude)	Only the text an AI feature needs for one request, with names replaced and contact details removed where we can. Not used to train models.
Vercel Inc. (AI Gateway)	Routes requests to the model provider	The same request in transit. Vercel states it deletes prompts and responses once a request completes.

We will give notice before adding or replacing an AI provider, as the Data Policy and the Data Processing Addendum say.

9. Reporting a problem with AI

If an AI feature gives a harmful, wrong or surprising answer, use **Report an issue** in the app (type: AI or the help assistant) or write to support@idealine.io. Do not include the sensitive text itself.

10. Changes

AI changes quickly. We may update this policy and will give notice of material changes as the Terms say.