

Responsible AI Policy

How Heyou's AI is built and constrained — Gemini via Vertex AI for Heyou's own AI processing, no training on your data, PII minimization before model calls, and no externally visible actions through users' accounts.

Version 1.1

Effective 10 May 2026

Last updated 6 October 2026

heyou.io/responsible-ai/

Heyou operates AI-assisted relationship intelligence. This policy explains what Heyou's AI does, what it cannot do, and how customers and individuals remain in control.

1. How Heyou's AI Works

Heyou combines:

1. A relationship graph built from Customer Data at the organizational level and User Data at the personal level;
2. Classifiers, scoring models, and path-ranking models trained or calibrated on relationship signals; and
3. Third-party large language models used to generate text, summaries, and action recommendations.

For personal-tier sources, Heyou processes metadata only, such as from, to, timestamp, interaction frequency, and similar non-content signals. Heyou does not process the content of a User's personal messages with other people. A question a User addresses to Heyou through a Heyou query channel is processed in order to answer it.

2. Model Inventory

Provider	Purpose in Heyou	Processing location	Training on Customer Data
Google Gemini via Vertex AI	Agent reasoning, message drafting, summarization, and recommendation support	Region-aligned where supported by provider, model, and endpoint configuration	Provider does not train foundation models on Customer Data submitted through enterprise controls
Anthropic Claude	Reasoning for questions a User asks through the Heyou Claude connector	The User's or Customer's own Claude environment	Governed by the User's or Customer's own agreement with Anthropic

The table above covers AI processing Heyou performs. Where a User queries Heyou from their own AI environment, such as the Heyou Claude connector, Heyou authenticates the User and returns the relationship data they are authorized to see; the model call itself runs in that environment under the User's or Customer's own arrangements with its provider. Questions asked through Slack, WhatsApp, or the Heyou app are served by Heyou's own orchestration on the providers Heyou operates.

Heyou maintains an internal evaluation process for each model and may add, replace, or retire providers in line with the subprocessor change-notice procedure.

3. What the Agents Do and Do Not Do

Heyou agents can:

- Suggest relationship paths and potential introducers;
- Draft outreach messages and replies for user review;
- Summarize relationships and prior interactions from authorized sources;
- Recommend next actions and timing;
- Answer a User's questions in Slack, WhatsApp, or Claude from data that User is already authorized to see, without reading those platforms' message or chat history or building any store from them, other than corrections a User asks Heyou to record.

Heyou agents do not, without explicit human-in-the-loop action:

- Send email or messages on a user's behalf;
- Perform externally visible engagement actions such as connection requests, comments, likes, follows, posts, or profile changes;
- Surface personal-tier User Data content or private signals to the Customer or other Authorized Users unless the individual User has affirmatively chosen to share specific data with the organization;
- Make decisions that produce legal or similarly significant effects on any individual;
- Access communication content beyond authorized integration scopes, and never the content of personal messages connected as User Data;
- Process special-category data, children's data, protected health information, payment card data, or protected employment decisions as part of the standard service.

4. Training-Data Posture

Heyou's AI posture distinguishes between third-party generative AI providers and Heyou's proprietary non-generative relationship-scoring models.

Third-Party Generative AI Providers

Customer Data sent to third-party generative AI providers is not used by those providers to train foundation models. Heyou uses Google Gemini via Vertex AI with enterprise controls designed to prevent provider-side model training on Customer Data. Retention protections and regional processing controls are configured where supported by the provider, model, and endpoint configuration.

PII Minimization Before Third-Party Calls

Where identifiers are not required for the task, Heyou minimizes direct identifiers before sending data to third-party generative AI providers. Names, emails, phone numbers, and similar identifiers may be replaced with placeholders, then re-associated inside Heyou after the model response returns.

This applies to model calls Heyou makes. Where a User queries Heyou from their own AI environment, such as the Heyou Claude connector, Heyou returns the relationship data that User is authorized to see and does not control the model call.

Heyou Proprietary Models

Heyou develops and improves proprietary non-generative classifiers, scoring models, path-ranking algorithms, and other relationship intelligence components. Heyou may use Customer Data and signals derived from it to improve these models in the course of delivering and enhancing the Services, subject to the DPA, applicable agreement, and applicable Order Form restrictions.

Heyou's default posture is customer-scoped model improvement: business-specific relationship intelligence is learned and applied within the customer or user context from which the data was authorized. Heyou's AI features are designed to use authorized tenant and user context, not unrestricted open-web access or uncontrolled external account access.

Google Workspace Data

Data received from Google Workspace APIs, including a User's connected Google Calendar, is subject to Google's Limited Use requirements, which are narrower than the posture described above and control where they conflict with it. Heyou does not use Google user data to develop, improve, or train generalized or non-personalized AI or machine-learning models. Calendar-derived learning is confined to the personalized model serving the individual User from whose account the data was authorized, and is not pooled into any model serving other users, other tenants, or Heyou's general-purpose models. The customer-scoped and tenant-scoped model improvement described above continues to apply to all other authorized data sources.

Safeguards include:

- No third-party generative AI foundation-model training on Customer Data.
- No generalized or non-personalized model training on Google Workspace data; calendar-derived learning stays within the authorizing User's personalized model.
- No individual cross-tenant leakage: one customer's personal data does not appear in another customer's outputs.
- No business-specific cross-customer reuse: one customer's relationship graph, messaging patterns, account strategy, outreach performance, sales motion, or GTM process is not used to generate business-specific recommendations for another customer.
- Tenant isolation: customer records are not retrieved in another customer's tenant.
- Aggregation and de-identification where feasible.
- Purpose limitation: model improvement is for relationship signal quality, path ranking, timing recommendations, safety, calibration, and service improvement, not for building or selling a contact database.

Any broader use of Customer Data for cross-customer model improvement requires support in the customer agreement, DPA, and applicable Order Form. Enterprise customers may address customer-specific model-improvement restrictions or opt-outs in the Order Form or DPA.

Statistical and Aggregated Insights

Heyou may derive aggregated, statistical insights from use of the Services, such as model calibration metrics or category-level benchmarks, provided such insights are de-identified and cannot reasonably be used to identify Customer, a Data Subject, or a natural person.

No Heyou-Owned Contact Database

Heyou does not combine data across customer tenants in its own database to build a contact directory, people-search product, or similar data product for sale or external use.

Future Changes

If Heyou materially changes its training-data practices in a way that expands customer-data use beyond the then-current agreement, Heyou will provide notice and obtain consent where required by contract or law.

5. Prompt Injection Controls

Relationship context can contain text written by third parties, such as a profile summary or a meeting title. Heyou treats that content as data, not as instructions: retrieved content cannot change an agent's instructions, expand its access, or cause it to act outside the requesting User's permissions. Agent tools are scoped to what the requesting User is already authorized to see, and query surfaces such as Slack, WhatsApp, and the Heyou Claude connector cannot act in third-party systems, so injected text has no path to an externally visible action.

6. Human Oversight

- Every agent-suggested action is reviewable by a human before it has external effect.
- Users can edit, approve, reject, or ignore AI-generated recommendations.
- Admins may configure organizational guardrails where supported.
- Action logs support auditability.

7. Transparency and Explainability

- AI-generated drafts and recommendations are presented as assistive outputs.
- Relationship scores may include top factors contributing to the score where supported.
- Heyou discloses approved model providers and purposes through its model inventory or enterprise review materials.
- Where a User interacts with Heyou through a query channel such as Slack, WhatsApp, or Claude, Heyou identifies itself as an automated agent rather than a person, and a route to human support remains available.

8. Fairness and Non-Discrimination

Heyou is not a tool for hiring, firing, credit, insurance, housing, immigration, compensation, performance evaluation, or other high-stakes automated decisioning. Customers must not use Heyou outputs for those

purposes. Heyou evaluates model behavior for quality, safety, and misuse risk as part of its AI governance program.

9. EU AI Act Alignment

Heyou's current intended use cases are relationship intelligence and human-led engagement assistance. Heyou is not intended for prohibited or high-risk AI uses. If Heyou's capabilities evolve into high-risk territory, Heyou will assess and implement applicable obligations before release.

10. Reporting AI Issues

Report suspected hallucinations, harmful outputs, or AI-related privacy concerns to dpo@heyoudo.io.