graine.ai

GRAINE TRUST CENTER · DOCUMENT

Responsible AI Statement

GR-PUB-08 · Version 1.3

Approved — effective 3 November 2026

DOCUMENT ID	VERSION	CLASSIFICATION
GR-PUB-08	1.3	Public

GRAINE AI · MAVRIX AI PRIVATE LIMITED

Document control

DOCUMENT ID	GR-PUB-08
VERSION	1.3
STATUS	Approved
OWNER	Founders
APPROVED BY	Rishabh Bhanot, Founder (top management), 3 Nov 2026
EFFECTIVE DATE	3 Nov 2026
LAST UPDATED	4 Nov 2026
NEXT REVIEW	By 3 Nov 2027 (12 months after approval), or on a significant change
CLASSIFICATION	Public
APPLIES TO	The AI features of the Graine AI business platform: hosted voice, chat and in-app agents, memory, transcript redaction and voice cloning

Contents

1. In short
2. What the platform does
3. Who is responsible
4. Telling people they are speaking with an AI
5. No training on customer data
6. Human oversight
7. Memory and redaction
8. Voice cloning and consent
9. Prohibited and restricted uses
10. Providers
11. Limits
12. Concerns and AI incidents
13. Where we stand
14. Contact
Related documents
Revision history

1. In short

This statement says how Graine AI approaches the responsible use of the AI in its platform, and what it does not yet do. Every statement is true as of 3 Nov 2026.

- Graine provides the platform. The customer decides what its agents say and whom they speak to.
- An agent can open by saying it is an AI and that the call may be recorded. This is built, off until switched on, and Being rolled out. Customers remain responsible for disclosure where the law requires it.
- Graine does not use customer data to train or fine-tune AI models.
- People stay in control: approvals, per-agent access, a call length cap, campaign pause, call stop, and transfer to a person where the call route supports it.
- Graine does not yet mark synthetic speech in a machine-readable way.
- Uses that affect people's legal position or similarly significant interests need human review and a specific written agreement. The platform is not a medical device.
- Graine uses ISO/IEC 42001 and the EU AI Act as reference frameworks. It is not certified to ISO/IEC 42001 and makes no claim of conformity with the EU AI Act.

2. What the platform does

Graine lets businesses build agents that hold conversations by phone, chat and inside mobile apps. An agent uses speech recognition, a language model and speech synthesis from providers named in the Sub-processor List. Intended uses are customer support, onboarding, reminders, scheduling and the qualification of enquiries.

3. Who is responsible

ROLE	RESPONSIBILITY
Founders (top management)	Accountable for how Graine builds and offers AI. They own this statement and decide what the platform will not be used for.
Information Security Lead	Keeps the AI controls, the provider list and this statement current.
Privacy & Grievance Officer	Answers privacy requests and complaints, including about the use of personal data by AI.
Customer	Decides each agent's purpose, words and audience; gives the notices and obtains the consents the law requires; reviews outputs; follows the Acceptable Use Policy.

4. Telling people they are speaking with an AI

An agent can open by saying that it is an AI and that the call may be recorded. An Admin can require this for the whole workspace. Both are built and Being rolled out: every rule is off until an Admin sets it, and the dashboard setting is not live yet. Until then, the customer's own greeting carries the notice. Art. 50(1) of the EU AI Act, which Graine uses as a reference, asks for people to be told they are dealing with an AI unless that is obvious; this notice is how Graine is building that in. The customer stays responsible for disclosure, recording notices and consent wherever the law requires them.

Graine does not yet mark synthetic speech in a machine-readable way, as Art. 50(2) of the EU AI Act describes. That is not implemented, and no date is set.

5. No training on customer data

Graine does not use customer data to train or fine-tune AI models. A custom voice is made only from recordings the customer gives for it. What AI and speech vendors do with data sent to them is set by their own terms. Graine has read each provider's published terms and is seeking written terms provider by provider. Some providers run model-improvement programmes of their own: several use content to improve their models by default unless the customer opts out, and one telephony provider, whose terms also cover the sign-in provider, uses customer content for its own purposes with no opt-out. The Security Questionnaire, which a customer can request under a confidentiality agreement, names them. Graine's decision (3 Nov 2026) is to opt out wherever a provider allows it: for one speech-to-text provider the platform-wide opt-out is built and is being switched on (Being rolled out), and the opt-outs in the other providers' settings are owner actions (Planned). Until that is done, ask Graine before choosing a provider for calls that matter.

6. Human oversight

- **Approvals.** Another person can be required to approve go-live, changes to a live agent and campaign launches. Nobody approves their own request.
- **Access.** Roles and per-agent access limit who can change which agent.
- **Stopping.** A campaign can be paused or cancelled and a call stopped. Every agent has a maximum call length; the default is 5 minutes.
- **Hand-over.** An agent can transfer a call to a person where the call route supports it.
- **Evaluations.** A customer can test an agent with simulated conversations before release. This is optional, and a workspace can require a passing test run before an agent goes live, from the dashboard or the API.

The customer decides when a person must check an agent's output before it is relied on.

7. Memory and redaction

- Memory is off until a customer turns it on. What is saved is redacted first, and is keyed by a one-way hash of the user id. One person's memory can be erased.
- When memory is on, it keeps names and health-related facts by default. This is a platform-wide setting, and customers should decide before enabling it.
- Memory expires after 365 days by default, which the customer can change. Phone-keyed memory ends 90 days after the last conversation.
- Transcript redaction is optional, per agent. It removes phone numbers, Aadhaar, PAN, card numbers and email addresses before storage. It does not find every personal detail. The in-app SDKs also mask those identifiers on the device.

8. Voice cloning and consent

A cloned voice is private to the member who made it. Nobody else can put it on an agent in Graine. A customer must have the speaker's consent before cloning a voice, and must not clone anyone's voice without it. A stored consent record for each clone is Being rolled out. A WhatsApp call goes out only when a consent from the person called is on record.

9. Prohibited and restricted uses

The Acceptable Use Policy governs. In summary:

- No deceiving people about whether they are speaking with an AI where disclosure is required, and no impersonating a person without their consent.
- No use that makes or materially drives decisions with legal or similarly significant effects on people, such as employment, credit, insurance or essential services, without human review and a specific written agreement with Graine first. These are the kinds of uses the EU AI Act's Annex III lists as high-risk.
- No health diagnosis or reliance on an agent in an emergency. The platform is not a medical device.
- No harassment, fraud, unlawful calling or practices the EU AI Act prohibits.

10. Providers

The language-model, speech and voice providers an agent uses are listed in the Sub-processor List, with what each handles and where. Graine gives workspace Admins 30 days' notice before it adds or replaces one, with a right to object. A provider a customer connects with its own account is the customer's processor.

11. Limits

Language models can state things that are wrong, and speech recognition can mishear, especially on poor lines or with background noise. An agent can therefore give a customer's caller a wrong answer. Customers should test their agents, review outputs, and give callers a route to a person.

12. Concerns and AI incidents

Harmful or false agent behaviour, undisclosed AI and misuse of a cloned voice are handled through Graine's incident process. Report a concern to founders@graine.ai. Graine tells an affected customer of a personal-data breach within 24 hours of confirming it. An incident register is In place (since 4 Nov 2026).

13. Where we stand

As of 3 Nov 2026:

IN PLACE	BEING ROLLED OUT	PLANNED
No training or fine-tuning on customer data; approvals; per-agent access; call length cap; memory controls; optional transcript redaction; on-device masking; member-only voice clones; WhatsApp consent record; the provider list on the Trust Center, with 30 days' notice of changes; and, since 4 Nov 2026, an optional rule that an agent goes live only after its latest test run passes, on the dashboard and through the API, an incident register, retention controls and data requests by person	Opting out of providers' model-improvement programmes platform-wide; AI-and-recording notice per agent and workspace; consent recorded for every dashboard voice clone; redaction by default for new agents; default calling hours (off until an Admin switches them on); an identity warning on memory	Machine-readable marking of synthetic speech; written AI impact assessments for new AI features and sensitive uses; per-agent switch to stop recording

Graine uses ISO/IEC 42001:2023 and the EU AI Act as reference frameworks to organise this work: a policy, named roles, impact assessment, controls across the life of an agent, third-party providers and responsible use. They are not claims of certification or conformity.

14. Contact

Questions, concerns or reports about the AI in Graine: founders@graine.ai.

Related documents

- Acceptable Use Policy
- Sub-processor List
- Shared Responsibility Matrix
- Security Overview
- Privacy Notice

Revision history

VERSION	DATE	CHANGE	APPROVED BY
1.0	2026-11-03	First issue, approved	Rishabh Bhanot, Founder
1.1	2026-11-03	Status of transcript encryption and the API go-live check	Status update; the approval of 3 Nov 2026 stands
1.2	2026-11-04	Vendor training defaults stated	Correction; the approval of 3 Nov 2026 stands
1.3	2026-11-04	Controls live after the 4 Nov 2026 deploy	Status update; the approval of 3 Nov 2026 stands