Artificial Intelligence (AI) Policy

HOPE INSIDE LIMITED

Document	AI Policy
Registered entity	HOPE INSIDE LIMITED
Company number	17228483 (registered in England and Wales)
Website	https://hopeinsidelimited.co.uk
Data Protection Lead	Ian Irving
Contact	DPO@hopeinsidelimited.co.uk
Effective from	19 May 2026
Version	1.0
Read alongside	Data Protection Policy

---

1. Why we have this policy

AI is central to what HOPE INSIDE LIMITED does. We publish an open-source multi-agent AI platform (AI Pandora's Box), use AI heavily in our publishing and media-production work (Pithos Media), and use AI to assist day-to-day group operations.

This policy explains:

- which AI tools we use, group-wide and per-brand;
- what data each tool sees, and what the tool's provider does with it;
- where a human reviews AI output before it reaches you;
- where AI is permitted to act without prior human approval, and where it is not;
- how you can ask for AI not to be used in your engagement with us.

This policy is read alongside our **Data Protection Policy** and applies to every member of the Hope Inside group team.

---

2. AI tools we use

Tool	Provider	Used for	Data the tool sees
Claude (Anthropic API)	Anthropic, PBC (US)	Drafting correspondence, classification, summarisation, code review, design + writing assistance across all group operations	The text submitted for processing

Tool	Provider	Used for	Data the tool sees
Gemini Live + Imagen + Veo (Google)	Google LLC	Real-time voice interaction; AI image + video generation for Pithos Media	Audio / text / image / video prompts submitted
ElevenLabs	ElevenLabs Inc. (US)	Narration, voice cloning (for our own brand voices only), audiobook production	Text submitted for synthesis; reference voice samples we own
Suno	Suno Inc. (US)	AI music generation for media projects	Text prompts
Microsoft Copilot (where contracted)	Microsoft Corporation	Productivity assistance inside our M365 tenant	Within our tenant only
GitHub Copilot	Microsoft / GitHub	Code-completion suggestions for our internal + open-source development	Source code in our development workspace
Local-model inference (Ollama, etc.)	Open-source models running on our hardware	High-volume classification, intent routing, draft-grade output that doesn't need a frontier model	All data stays on our hardware -- never leaves the building

We do not use AI tools for:

- Facial recognition, voice identification, or any biometric matching of individuals without consent;
- Automated decisions about you that have legal or similarly significant effect (Art. 22 UK GDPR);
- Scraping of personal data for training a public model without explicit written consent of the person concerned;
- Producing realistic synthetic media of identifiable real people (deepfakes) outside of consent-based brand work.

---

3. Human-in-the-loop -- what that actually means

Every external-facing output -- a customer email, a published article, a released audio or video clip, a code release -- is reviewed and signed off by a human before publication.

Action	Approval gate
Email to a customer or contributor	A human reads + sends. AI drafts only.
Published article / blog / video script	A human edits + publishes.
Audiobook chapter / narration delivered to a customer or platform	A human listens to the full take + signs off.
AI-generated music or video published	A human reviews the final cut + brand-alignment + publishes.
Software release (Pandora's Box, etc.)	A human authors the release notes + tag. AI may suggest changelog wording.

Action	Approval gate
Customer-support classification / routing	A human reviews on first response; prior-approved routing is allowed for the same correspondent.
Marketing campaign send	A human reviews + sends.

Internal-only AI actions (e.g. scheduled internal classification jobs in our admin tooling) are logged, monitored, and reviewable on request.

---

4. Data handling for AI calls

- We use each provider's **zero-retention** API path where one is offered.
- We **do not** opt in to providers using your data to train their public models.
- We **redact** account credentials, payment details, and unnecessary personal data before sending.
- We **log** every AI call on our side (time, tool, purpose, byte counts).
- AI output referencing your data is treated as your data, with the same retention rules as in the Data Protection Policy.

Each provider relationship is governed by a written agreement that includes the UK International Data Transfer Agreement or equivalent safeguard, where transfers leave the UK.

---

5. Open-source contributors -- specific notes

If you contribute to one of our open-source projects (notably AI Pandora's Box):

- The contributor-licence agreement governs how we use your contribution.
- We may run AI tools (Claude, etc.) against the public repository for code review, summarisation, and release-notes drafting. Your contribution is public-record, so this is not a personal-data disclosure beyond what you have already published.
- We will not use your private correspondence with us (issue triage discussions held privately, email threads) to train any public AI model.

---

6. Your choices

You can ask, in writing, that:

- AI tools should not be used in your engagement with us;
- AI tools may be used for drafting but not for classification or routing of your correspondence;
- We disclose, for any given output, which AI tool (if any) was involved in producing it.

Send your request to **DPO@hopeinsidelimited.co.uk**. We will confirm within 5 working days.

---

7. AI errors and accountability

The human who released the output is accountable, not the AI. We document errors, identify root causes, and adjust our use of the tool. Where the error caused harm or loss, our standard service-agreement terms apply.

We will not hide behind "the AI got it wrong" as a defence.

---

8. AI literacy and staff training

Anyone in our team who uses an AI tool on customer or contributor-facing work has been trained on:

- the purposes for which each tool is permitted;
- the data-handling rules in Section 4;
- the human-in-the-loop expectations in Section 3;
- how to recognise and report a suspected AI-generated hallucination, bias, or data leak.

Training is refreshed at least annually and on each material change to our AI tool list.

---

9. Open-source ecosystem note (AI Pandora's Box)

AI Pandora's Box is published under Apache 2.0. Anyone can self-host and use it. We do not have visibility into how third-party operators use the platform in their own environments; the platform documentation makes clear that those operators are themselves the data controllers for any personal data they put through their installation. We are not a data controller for personal data processed inside someone else's self-hosted Pandora's Box install.

---

10. Changes to this policy

Updated when we add, remove, or materially change an AI tool, or when applicable UK law or ICO guidance materially changes. The current version is dated at the top.

HOPE INSIDE LIMITED · Company number 17228483 · Registered in England and Wales