An AI use policy *a board can adopt.*

This is the policy LUR Growth works under on every engagement, and the one we hand to your board to adopt as its own. Read it in ten minutes, adopt it in one meeting, apply it on a Tuesday. Replace the bracketed words with your organization's names and it is yours.

1. Purpose

[Organization] uses AI tools to reduce the manual work that keeps staff from the people they serve. This policy sets the rules for how those tools may be used, what information may reach them, who reviews their output, and who is accountable. It exists to protect the people in our records, the trust of our funders, and the judgment of our staff.

2. Scope

Applies to every employee, board member, volunteer, and contractor who uses an AI tool with [Organization]'s information, whether the tool is one the organization provides or one a person brings. Covers drafting, summarizing, research, data handling, and any automated workflow that uses an AI model.

3. Approved uses

- Drafting first versions of reports, acknowledgments, board packets, newsletters, and grant narratives from records the organization already holds.
- Summarizing meeting notes, documents, and public research.
- Reformatting and reconciling data the organization already owns.
- Automated workflows that a named person on staff reviews before any output reaches a donor, client, family, funder, or trustee.

4. Uses that need approval first

- Any workflow that sends information about an identifiable person to an AI tool for the first time.
- Any use that touches client, patient, student, or family case records.
- Any new AI vendor, or a change to an existing vendor's terms.
- Any output that will be published or filed without a person editing it.

Approval comes from [the executive director or designee], in writing, and is recorded in the workflow's documentation.

5. Prohibited uses

- Entering passwords, bank details, Social Security numbers, health diagnoses, or immigration status into any AI tool. For minors, nothing beyond the fields a school's or program's own enrollment record holds and the specific workflow needs.
- Using an AI tool to decide a person's eligibility, benefits, employment, or care without a person making the final call.
- Using a personal or free consumer AI account for organizational information.
- Presenting AI-generated text, images, or numbers as a person's own statement or as verified fact without review.

6. Data rules

Information falls into three tiers. The tier decides what may reach an AI tool.

- **Never sent.** Credentials, financial account numbers, government identifiers, health and legal details about identifiable people, anything a person has asked us to keep confidential, and any full export of a database.
- **Sent only as the fields a draft needs.** Names, giving history, program participation, and contact details, and only the fields the specific workflow uses. Never the whole record.
- **Sent freely.** Public information, the organization's own published material, and templates.

The database of record stays the organization's. AI tools read from it and draft against it. They do not become it.

7. Vendors and terms

Drafting and automation run on Claude, by Anthropic, under commercial terms that keep the organization's inputs out of model training. Every other vendor in a workflow (the database, the automation platform, the email system) is named in that workflow's documentation, with the account owner and the terms it runs under. All accounts are owned by [Organization].

Contractors build inside them and do not hold organizational data in accounts of their own once the work is handed over.

8. Human review

Every workflow names one person on staff as its reviewer. Nothing an AI tool drafts reaches a donor, client, family, funder, trustee, or the public until that person has read it and approved it. The reviewer may edit, hold, or discard any output, and is never asked to approve faster than they can read.

9. Accuracy, bias, and voice

AI tools make confident mistakes. Reviewers check every number, name, date, and dollar figure against the source record before approval. Any output that describes a person or a group is read for bias and for language the organization would not use in its own voice. When in doubt, the reviewer rewrites it or sends it back.

10. Ownership and documentation

Every workflow is documented as it is built: what it does, what data it touches, who reviews it, which vendors it uses, and how to turn it off. Documentation, admin rights, and training recordings belong to [Organization]. No workflow depends on knowledge held only by one person, staff or contractor.

11. Transparency

When someone asks whether a piece of communication was drafted with an AI tool, we say so plainly. Funders and partners who ask how AI is used in our reporting receive this policy.

12. Reporting a concern

Anyone who sees an AI tool used outside this policy, or an output that is wrong or harmful, reports it to [the executive director or designee] the same day. The workflow is paused until the concern is resolved. Reporting a concern in good faith is never held against the person who raised it.

13. Review

The board reviews this policy once a year, and sooner if a vendor changes its terms, a new kind of data enters a workflow, or an incident is reported. Each review is recorded in the board minutes.

Adoption

Resolved, that the board of [Organization] adopts this AI Use Policy, version 1.0, effective [date], and directs [the executive director] to apply it to every current and future AI workflow and to report on its use at least once a year.

Chair: _____ Date: _____

How LUR Growth applies it. Built in accounts you own. Only the fields a draft needs are sent. A named person reviews every output. Everything is handed over. lurgrowth.com/ai-use-policy · hello@lurgrowth.com