AI Policy: A Leader's Guide for Nonprofits

Five decisions, a working group to lead the change, a green, yellow, red sort for everyday work, and an eight-week path to adoption.

Make the decisions before you open a template

If your staff use AI, they may be paying for it themselves, using a free version, or using what arrived inside software you already own. An AI policy is how leadership decides what is acceptable in each case.

Free templates are easy to find. Each one asks the same questions: who approves tools, which information can go into which tool, and what needs review. Those are decisions for leadership, and a template cannot make them for you. Make those decisions first. Writing the policy takes much less time once they are settled.

Five decisions leadership makes first

1. **Who owns it and how tools are approved.** Decide who adopts the policy and who owns it. For example, the board adopts it, the executive director owns it, and a small AI working group drafts it and leads the rollout. Then decide how a new tool is approved. For example, your IT lead or provider reviews the tool for security first, the working group tests it, and the executive director gives final approval. The executive director may ask counsel for an opinion before deciding.
2. **Which tools, on which accounts.** Name the approved tools, and put them on accounts the organization pays for and controls. The organization has no control over a personal account. It cannot see what was entered, set how long conversations are kept, or close the account when someone leaves. Before a tool is approved, read the vendor's terms for who owns what you enter, where it is stored, and what happens to it when you cancel.
3. **Which information goes in which tool.** List the kinds of sensitive information you hold, in plain words: client and case records, health information, anything about a minor, donor payment details, Social Security and other ID numbers, personnel files, and anything under a confidentiality agreement or a funder's data terms. Then decide which tool, if any, is approved for each kind. An assistant built into your own Microsoft 365 or Google Workspace tenant, such as Microsoft 365 Copilot or Gemini in Google Workspace, works under the organization's agreement and follows the file permissions you set, so leadership can approve it for records that should never go into a personal chatbot account. Where a health privacy law, card payment rules, or a funder's terms apply, confirm that the provider's agreement covers the tool first. Staff should not put anything on the list into a tool that is not approved for it.
4. **What needs human review.** AI can prepare a draft, but a person reviews it and makes the decision. State which work must be reviewed before it leaves the organization: anything going to a funder, a donor, a client, or the public. The person who sends it is responsible for every fact and figure in it. Reviewers also check that text and images do not copy someone else's work, and that a recommendation about a person, such as eligibility or hiring, is fair before anyone acts on it. Decide when you will tell a donor, client, or funder that AI was used.
5. **What happens when something goes wrong.** Name the person staff should tell, and ask them to report the same day. Make it clear that reporting an honest mistake will not be held against them.

The AI working group

Someone has to lead the change after the decisions are made. Form a small AI working group and give it that job. Include someone from programs, development, finance or HR, and whoever runs your technology. Include at least one person who already uses AI and one who has doubts about it.

Before the group starts its work, leadership and the group agree on a goal for AI in the organization: what it should help with, such as less staff time spent on routine writing and reporting, and what it will not be used for. Write the goal down. It tells the working group what it is working toward.

The working group drafts the policy and the list of green, yellow, and red uses. It tests a new tool before it is approved, brings staff questions and concerns to leadership, and helps train their own teams. It keeps meeting after the policy is adopted, to review new tool requests and the problems staff run into. The group makes recommendations, and the policy owner makes the decisions.

Sort everyday uses into green, yellow, and red

Give staff examples along with the rules. Sort your own everyday uses into three groups, using the fillable worksheet on the last page, and post the list where people can see it.

GROUP	WHAT IT MEANS	NONPROFIT EXAMPLES
● **Green**	Permitted on an approved tool	Drafting a newsletter from public information. Summarizing a published report or a funder's public guidelines. Brainstorming event themes. Editing your own draft when it contains nothing confidential.
● **Yellow**	Requires approval first, and review of the result	Grant proposals and reports. Donor messages that use names or giving history. Client or personnel records, on the tool approved for them. Recording and transcribing meetings. Sorting resumes or volunteer applications. Connecting an AI tool to email, files, or the donor database. AI-made images of people.
● **Red**	Not permitted	Client, health, personnel, or other sensitive records in a tool not approved for that data. Organization work on a personal account. Letting AI make the final decision on eligibility, hiring, or services. Publishing AI output that no one has reviewed.

Some yellow items have legal requirements, such as recording a meeting or screening applicants. Ask counsel what the law requires for those.

Six questions board members can ask

Any question leadership cannot answer is the place to start.

1. Who, by name, owns AI use in this organization?
2. Which tools are approved, and are they on accounts we control?
3. What sensitive information do we hold, and which tool, if any, is approved for each kind?
4. What can the assistant inside our email and file system read today?
5. When were staff and volunteers last trained on this?
6. If something goes wrong tomorrow, who is called first?

What a policy cannot do alone

A policy says what the organization intends. Four things have to be in place for it to work.

- **The files staff can already open.** AI assistants built into Microsoft 365 and Google Workspace can draw on the supported files the person asking has permission to open, unless added security controls block them. That can include a salary spreadsheet shared with everyone years ago, a board folder that was never restricted, or a client list on an old team site. The assistant can bring up any of them when someone asks a question. Review file sharing before the assistant is turned on, and again before it is approved for sensitive records.
- **The settings behind the policy.** A policy can state that your data is not used to train AI models. Whether that is true depends on the plan you are on and how it is set up. On many consumer AI accounts, what people type can be used to improve the product, depending on a setting each person controls. Business plans from Microsoft, Google, OpenAI, and Anthropic state that customer content is not used to train their models by default. Confirm the plan, the training setting, how long conversations are kept, and that a person's access ends on their last day.
- **Training for staff and volunteers.** Run a short session built on your own green, yellow, red list. Volunteers and board members who handle the organization's information need it as well. Give them a one-page version and include it in onboarding.
- **An incident plan.** Decide in advance what counts as an AI incident and who is notified. Plan for at least these four: confidential information pasted into a tool that was not approved for it, an AI-written error sent to a funder or published, a stolen sign-in to a tool that holds months of conversations, and a cloned voice or video of a leader asking for a payment. For each one, name who is called: the executive director, your IT provider, counsel, and your insurer.

An eight-week path

Plan on about eight weeks. A larger organization may need more time.

WEEKS	WHAT HAPPENS
1 and 2	**Set the goal and form the group.** Leadership names the policy owner and forms the AI working group. Together they agree on the goal for AI in the organization and write it down.
3 and 4	**Find out what is in use.** The working group asks staff which AI tools they use and what they use them for. Tell staff no one will be in trouble for an honest answer, and hold to that.
5 and 6	**Decide and draft.** Leadership makes the five decisions, with the working group's recommendations. Sort your ten most common uses into green, yellow, and red. Fill in a template with what you decided.
7 and 8	**Review, adopt, and train.** Have counsel review the draft. Check your grant agreements, funder terms, and cyber insurance application for AI or data-handling conditions. The board adopts the policy. Train staff and volunteers, and confirm the settings on every approved tool. Schedule the next review in six months.

Texas laws to raise with counsel

These Texas laws bear on the decisions in this guide. This is a short summary and not legal advice. Ask counsel how each one applies to your organization.

- **Texas Responsible Artificial Intelligence Governance Act** (Business and Commerce Code, Chapters 551 and 552). In effect since January 1, 2026. It applies to anyone who promotes, advertises, or conducts business in Texas, produces a product or service used by Texas residents, or develops or deploys an AI system in the state. It has no general exemption for nonprofits. Among other things, it prohibits developing or deploying an AI system with the intent to unlawfully discriminate against a protected class. The attorney general enforces it.
- **AI in health care.** When an AI system is used in a health care service or treatment given by a licensed, registered, or certified provider, the provider must disclose that to the patient or the patient's personal representative no later than the date the service or treatment is first provided, or as soon as reasonably possible in an emergency (Business and Commerce Code, Section 552.051). A health care practitioner who uses AI for diagnostic purposes, including recommendations on a diagnosis or course of treatment, must review all records created with AI in a manner consistent with Texas Medical Board standards and must disclose that use to patients (Health and Safety Code, Section 183.005).
- **Health information** (Health and Safety Code, Chapter 181). The Texas medical privacy law can reach organizations that the federal HIPAA rules do not, because it covers any person who comes into possession of protected health information. It has exemptions, including one under Section 181.053 for certain nonprofit agencies that pay for health care or prescription drugs for indigent people. A covered organization must train the employees who handle that information.
- **Breach notice** (Business and Commerce Code, Section 521.053). It applies to a person who conducts business in Texas and owns or licenses computerized data that includes sensitive personal information. When that information is acquired, or reasonably believed to have been acquired, by an unauthorized person, the people affected must be told without unreasonable delay, and no later than 60 days after the organization determines that the breach occurred. If at least 250 Texas residents are affected, the attorney general must be told as soon as practicable, and no later than 30 days after that date. Ask counsel about the exceptions.
- **Recording** (Penal Code, Section 16.02). Recording a conversation can be a crime in Texas. The law provides a defense when the person recording is a party to the conversation, or when one of the parties consented beforehand, unless the recording is made in order to commit an unlawful act. Some states require everyone's consent, which matters when a call includes people outside Texas.
- **Consumer privacy** (Business and Commerce Code, Chapter 541). The Texas Data Privacy and Security Act does not apply to nonprofit organizations.

The official text of each law is at statutes.capitol.texas.gov.

Resources

These resources are free to read, and some link to paid courses. Use a template once the five decisions are made, and have counsel review the result before the board adopts it.

Policy templates

- **Fast Forward, Nonprofit AI Policy Builder.** Answers a set of questions and produces a draft. ffwd.org/nonprofit-ai-policy-builder
- **Community IT, Acceptable Use of AI Tools in the Nonprofit Workplace.** A short template, last revised May 2026, free with a short form. communityit.com
- **TechSoup, AI for Nonprofits.** Guides, courses, and a generative AI usage policy template in Word format. page.techsoup.org/ai-for-nonprofits

Background reading

- **NTEN, AI for Nonprofits Resource Hub.** Articles, courses, and a policy template for nonprofit staff. nten.org
- **Fundraising.AI, Framework for Responsible AI.** Ten principles for AI in fundraising and donor trust. fundraising.ai/framework
- **NIST AI Risk Management Framework.** A voluntary federal framework for managing AI risk. nist.gov

Where your data goes, by provider

- Each page covers specific plans, so check it against the plan you are on: Microsoft 365 Copilot, Gemini in Google Workspace, ChatGPT business plans, and Claude.

Security

- **Mount Xion, The Cyber Risks Facing Texas Nonprofits.** Our report on what Texas nonprofits are losing to fraud and breaches, with a free PDF. mountxiontechnologies.com/insights

This guide is general information for nonprofit leaders. It is not legal advice.

Where Mount Xion fits

Mount Xion Technologies is an SBA-certified Service-Disabled Veteran-Owned Small Business in San Antonio, Texas. Our AI adoption and governance work covers what a policy cannot do alone: a review of what your AI tools can read today, the settings behind the policy, an acceptable-use policy written around your tools and your work, and staff training on real examples.

The first conversation is a free 30-minute discovery call.

Book a call	mountxiontechnologies.com/book
The service	mountxiontechnologies.com/ai-adoption
Email	info@mountxiontechnologies.com
Phone	(210) 729-0079

Worksheet

Type into this page on screen, or print it and fill it in by hand.

Our goal for AI

Policy owner

Working group members

Our own list

Write down the ten things your staff most often use AI for, or want to, and mark each one green, yellow, or red.

HOW WE USE AI, OR WANT TO	GREEN	YELLOW	RED
	☐	☐	☐
	☐	☐	☐
	☐	☐	☐
	☐	☐	☐
	☐	☐	☐
	☐	☐	☐
	☐	☐	☐
	☐	☐	☐
	☐	☐	☐
	☐	☐	☐