

A GUIDE FOR NONPROFITS, CHURCHES AND MINISTRIES

Before AI Speaks for Your Organization

Why every serious organization needs a written AI policy, what can happen without one, and how to have yours finished and approved this week.

This guide

Why it matters, what's at stake, and how to do it

AI Use Policy Template

Word document to fill in and adopt

Plain-text version

Ready to paste into your AI tool

AI is already inside your organization.

Someone on your staff, your board or your volunteer team has already pasted something into a chatbot this month. A newsletter draft. A grant paragraph. Maybe a story about someone you serve. The question isn't whether your organization uses AI. It's whether anyone has decided *how*.

92%

of nonprofits now use AI in some form.

Virtuous, 2026 Nonprofit AI Adoption Report

47%

have no AI policy at all. No rules, no limits, no one accountable.

Virtuous, 2026

7%

use it well enough to see real impact on their mission.

Virtuous, 2026

Those numbers describe a sector moving fast without guardrails. Most AI use in small organizations happens quietly, one person at a time, on personal accounts, with no shared idea of what is allowed. Leaders tend to be enthusiastic; staff tend to be uneasy. A 2026 NTEN and Bridgespan survey found executives far more likely than their staff to see AI as a help, and most organizations still had no written rules or plan for it.

That gap is where the risk lives. When there are no rules, each person makes up their own. Some will be too careful to use a genuinely useful tool. Others will paste in things that should never leave your building.

An AI policy doesn't slow anyone down. It is the line that lets people move quickly without worrying they're about to make a serious mistake.

When AI gets it wrong, the organization pays for it.

These are real, public cases. None of them involved bad intentions. Each one involved AI being used without clear limits.

NONPROFIT · 2023

A helpline chatbot gave harmful advice

A national eating-disorder nonprofit took its support chatbot offline after users reported it giving weight-loss and calorie-counting advice: the exact opposite of its mission, delivered to the people most at risk.

Never let AI speak to vulnerable people without a person in the loop.

AIRLINE · 2024

"The chatbot said so" wasn't a defense

A Canadian tribunal ordered Air Canada to honor a refund policy its website chatbot had invented. The company argued the bot was responsible for its own words. The tribunal disagreed: the organization owns what its AI says.

Anything AI says in your name, you said.

LAW FIRM · 2023

Invented sources in an official filing

Two New York lawyers were fined after filing a brief that cited court cases ChatGPT had made up. The cases sounded real. Nobody checked.

A grant report or a funder letter is an official document too. Every fact must be verified.

CORPORATION · 2023

Confidential material pasted into a chatbot

Samsung restricted staff use of generative AI after employees pasted confidential internal material into ChatGPT, where the company no longer controlled it.

What goes into an AI tool may not stay private. Decide in advance what never goes in.

If a global airline, a law firm and a tech giant can get caught out, a small team working late on a Thursday can too. The difference is that a small organization has much less room to recover.

The quiet ways it goes wrong in a small organization.

Most AI damage never makes the news. It shows up as a lost donor, an uncomfortable call from a funder, or a person you serve who stops trusting you.

- 1

A real story with a real name. A volunteer pastes a client's case notes into a free chatbot to "make it sound nicer" for the newsletter. That person's history, maybe their addiction, record or diagnosis, is now somewhere you don't control.
- 2

A number nobody can back up. AI fills a gap in a grant report with a statistic that sounds right. The funder asks for the source. There isn't one.
- 3

A photo that isn't real. An AI-generated picture of "the kids in our program" ends up in an appeal. A donor recognizes it as AI. Now they wonder what else isn't real.
- 4

A voice that isn't yours. Every email starts sounding like every other organization's. Donors who gave because you felt personal stop opening.
- 5

Words that wound. AI describes the people you serve in language you would never use. It goes out because nobody was assigned to read it first.
- 6

Work that walks out the door. A staff member builds everything in a personal AI account, then leaves. Your prompts, drafts and history leave with them.

The people you serve

Their privacy and dignity are the first thing at risk.

Your donors

9 in 10 nonprofit leaders say donor trust is already harder to earn.

Your funders

An unverifiable claim can cost a grant, and a reputation with it.

Your board

Overseeing risk is part of a board's duty. No rules means no oversight.

Your staff

Without clear rules, people either freeze up or guess.

One document, four protections.

The template in this kit brings together four things most organizations either don't have or keep scattered across people's heads. Together, they tell both your people *and* your AI tools exactly how to act on your behalf.

1–2

Who you are and how you sound

Your exact mission, who you serve, your approved facts, and your brand style: voice, words you use and words you never use, colors, fonts, who signs your letters.

PROTECTS

Your voice, your accuracy, and how the people you serve are described.

3

The rules for using AI

Which tools are approved, what AI may help with, what needs extra sign-off, who reviews everything before it goes out, and what to do when something goes wrong.

PROTECTS

Quality and accountability. Every piece of work has a named human responsible for it.

4

Governance from your by-laws

Who may speak for you, sign for you, spend for you and submit grants for you, taken straight from the by-laws you already have, so AI-assisted work follows the same chain of authority as everything else.

PROTECTS

Your board's authority, and your legal and financial lines of responsibility.

5

What AI must **never** do

The red lines that don't change for any deadline: no personal information in, no invented facts or stories, no fake photos, no decisions about people, no advice on your behalf, nothing published without human review. A starter list is included. You add your own.

PROTECTS

Everyone. This is the part that prevents the cases on the previous pages.

Numbers match the parts of the template. It also includes Part 6 (instructions written for the AI itself) and Part 7 (adoption and sign-off).

It protects you. It also makes AI actually useful.

The organizations getting real value from AI aren't the ones using it most. They're the ones that told it clearly who they are and what it may and may not do.

You sound like yourself

Once your voice, terms and approved facts are written down, every draft starts much closer to finished. Less rewriting, fewer "this doesn't sound like us" moments.

You can safely hand work to others

Volunteers, interns and new staff can use AI for real work on day one, because the rules are on paper rather than in one person's head.

You show funders you're well run

Funders, insurers and auditors increasingly ask how organizations handle AI and data. A board-adopted policy is a clear, simple answer.

The people you serve stay protected

Their names, histories and images stay out of tools you don't control. That is a promise of dignity as much as a privacy rule.

Your board can do its job

Oversight needs something to oversee. A written policy, adopted and reviewed yearly, turns a vague risk into a governed one.

Your team stops guessing

Staff who were nervous about AI get permission and limits. Staff who were too relaxed about it get guardrails. Everyone works from the same page.

Legitimate organizations write down how they handle money, people and data. AI now touches all three.

Simpler than it sounds.

You don't need a lawyer to start, a consultant, or a committee that meets for six months. You need the template, your by-laws, and about two focused hours.

-
- | | | |
|---|-------------------------------------|---------------|
| 1 | Gather what you already have | 15 min |
| Your mission statement, your by-laws, a letter or post you're proud of, your brand colors and fonts, and a list of the AI tools people are already using. | | |
| 2 | Fill in the template | 45–60 min |
| Most of it is short answers and checkboxes. Start with Part 5, the NEVER list. If you only finish one section today, make it that one. | | |
| 3 | Adopt it officially | 1 agenda item |
| Bring it to your board or leadership. Vote on it and record it in the minutes. That is what turns a document into policy. | | |
| 4 | Load it into your AI tool | 10 min |
| Upload the finished file (or clear photos of each page) to your AI tool. Put it in a Project or Gem so it applies to every chat (see the next page). | | |
| 5 | Share it, sign it, review it | Ongoing |
| Everyone who uses AI for your work reads it and signs Part 7. Put a review date on the calendar for one year from now. | | |
-

Total time to be covered

About two hours

Same request. Two very different results.

THE REQUEST "Write a thank-you email to a first-time donor."

WITHOUT A POLICY

Dear Valued Donor,

Thank you for your generous contribution!
Thanks to supporters like you, we've helped
over 5,000 families this year and transformed
countless lives.

Your donation is 100% tax-deductible.
Together, we're making a difference!

With gratitude,
The Team

- ✗ Invented a statistic
- ✗ Made a tax promise
- ✗ Sounds like everyone else
- ✗ No one named as sender

WITH YOUR POLICY LOADED

Dear [DONOR FIRST NAME],

Your first gift arrived this week, and I wanted
you to hear from me directly. Because of you,
someone walking through our doors this
month will find a hot meal and a person who
knows their name.

This year we served [VERIFY: families served,
2026] families. Soon I'll send you a short story
about one of them.

Gratefully,
Pastor [NAME]

- ✓ Flags the number to verify
- ✓ No tax promise (Part 5, rule 13)
- ✓ Your voice and your sender
- ✓ Placeholders, no personal data

Giving it to your AI: as easy as sending a photo

Upload the finished policy (the Word file, a PDF, or clear photos of every page) and type: **"This is our AI policy. Follow it for everything you write for us."** The AI reads it and follows it. A file is better than photos: a blurry or missing page could drop one of your NEVER rules.

Make it stick. Uploaded in a regular chat, the policy only lasts for that conversation. Upload it once inside a Project or Gem and it applies to every chat there:

ChatGPT

Make a Project and add the policy to its files or instructions.

Claude

Make a Project and add the policy to its project knowledge.

Gemini

Make a Gem and upload the policy as its knowledge.

Microsoft Copilot

Add the policy to a notebook or agent your team shares.

You're covered when...

-
- ☐ Your mission, approved facts and brand style are written down in one place.
 - ☐ You've listed which AI tools are approved, and they're organization accounts where possible.
 - ☐ Every piece of AI-assisted work has a named human reviewer before it goes out.
 - ☐ Your by-laws' lines of authority (who speaks, signs, spends and submits) are reflected in the policy.
 - ☐ Your NEVER list is complete, and nobody is unclear about it.
 - ☐ Your board or leadership has formally adopted it, and it's in the minutes.
 - ☐ The policy is loaded into every AI tool your team uses.
 - ☐ Everyone who uses AI for your work has read and signed it.
 - ☐ A review date is on the calendar.
-

The people you serve trust you with their stories. Your donors trust you with their money. Your community trusts that what you say is true. A written AI policy is how you keep all three, now that a machine can write in your name.

This guide and template are not legal advice. Have your board review your policy. Where your work involves confidential, health, legal or children's information, ask a qualified attorney to review it as well.

Sources. Virtuous, *2026 Nonprofit AI Adoption Report* (survey of 346 nonprofits). NTEN and Bridgespan nonprofit AI survey, 2026, via *Chronicle of Philanthropy*. Momentive Software survey of 500 nonprofit executives, 2026. National Eating Disorders Association chatbot "Tessa," suspended June 2023. *Moffatt v. Air Canada*, BC Civil Resolution Tribunal, February 2024. *Mata v. Avianca*, U.S. District Court, S.D.N.Y., June 2023. Samsung generative-AI restriction, reported May 2023.