

---

A PRACTICAL GUIDE FOR MISSION-DRIVEN ORGANIZATIONS

# AI for Nonprofits

*From first question to first tool*

---

Artificial intelligence isn't just for tech companies. Nonprofits are using it right now to write grant reports faster, answer donor questions at midnight, translate materials for multilingual communities, and stretch small staff teams further than ever before. This guide covers what to consider, what to build before you launch, and how to move forward without creating problems you'll spend months cleaning up.

## WHAT AI DOES WELL

Drafting, summarizing, translating, answering repetitive questions, analyzing data, and generating first drafts of almost anything written.

## WHERE IT SAVES TIME

Grant writing, donor communications, volunteer onboarding materials, social media content, intake forms, and program reporting.

## WHAT IT CAN'T REPLACE

Human relationships, community trust, ethical judgment, and the lived experience your staff and clients bring every day.

## ON ETHICAL AI

Ethical AI means choosing tools that are transparent about how they work, protecting client and donor data, being honest with your community when AI is involved, and maintaining human oversight over every decision that affects the people you serve. AI should expand your capacity, not replace your accountability.

# Contents

---

## QUICK REFERENCE

---

- » Quick reference summary
- 

## PART 1: BEFORE YOU DECIDE

---

- 01 Is AI right for your organization right now?
  - 02 What to do before you come back to this
  - 03 When to keep AI out entirely
- 

## PART 2: GETTING STARTED

---

- 04 The process at a glance
  - 05 Planning categories
  - 06 Tools to consider
  - 07 Questions to ask every vendor
  - 08 Skills, hiring, and training
- 

## PART 3: DOING IT RESPONSIBLY

---

- 09 What human oversight actually looks like
  - 10 How to know if it's working
  - 11 What to tell your board, funders, and community
  - 12 When AI faces your clients directly
  - 13 Procurement and vendor contracts
  - 14 AI use policy template
- 

## PART 4: WHEN YOU NEED IT

---

- A Terms you'll encounter

---

- B What goes wrong in the first year

---

- C Physical and data security

---

- D Local models vs. cloud tools

---

- E Minnesota-specific resources

---

- F AI and grant writing

---

# Quick reference summary

A condensed version of this guide for sharing with colleagues, presenting to a board, or keeping on hand during implementation. Everything here is covered in more depth in the sections that follow.

## 5 questions before you start

- 1 What specific task is costing us the most time. Is AI actually suited to it?
- 2 Who will lead this internally, and do they have the time?
- 3 Does the tool we're considering have a clear data privacy policy. Does it meet our standards?
- 4 Do we serve populations where AI errors or AI disclosure would have serious consequences?
- 5 What does success look like at 90 days. How will we measure it?

## 3 things AI cannot replace

### Community trust

The relationships your staff have built with the people you serve are not automatable. AI can help you communicate more efficiently. It cannot build or repair trust.

### Ethical judgment

AI applies patterns. It does not weigh competing values, consider context, or hold your organization accountable to its mission. That requires a human being.

### Accountability

When something goes wrong (and at some point it will), a human being needs to be able to say "I reviewed this, I approved it, and I take responsibility for it." AI cannot be accountable. Your staff can.

## Who needs to be at the table

**Executive Director:** sets the mandate, approves the policy, reports to the board

**Internal AI lead:** owns implementation, training, and ongoing oversight

**Program staff:** know where the time is lost and where the risks are highest

**Development staff:** often the first to adopt, need clear funder communication guidance

**At least one board member:** governance, not operations, but informed and engaged

## The short version of this entire guide

Start with one use case, not five. Pick the task costing your team the most time. Choose a tool with clear data privacy terms. Write a one-page policy before you launch. Train staff with hands-on time, not a handout. Review every AI output before it leaves the building. Tell your board, your funders, and your community what you're doing and why. Measure what changes at 30, 60, and 90 days. Adjust before you expand.

AI will not save your organization. It will give your team more time to do the work that will. That's a meaningful difference. It's enough of a reason to do this carefully.

### THE ONE THING TO REMEMBER

*AI should expand your capacity, not replace your accountability. Every decision that affects a person should have a person behind it.*

### PART 1

## Before you decide

Two sections that belong together: whether AI is right for your organization now, and what to do if the timing isn't right. Work through these before anything else.

# Is AI right for your organization right now?

Adding AI is an investment: in time, training, and ongoing attention. Answer these questions honestly before moving forward.

## Signals you may be ready

- Staff spend significant time on repetitive writing tasks
- You have at least one person willing to lead the effort
- Leadership is open to learning alongside staff
- You have clear data privacy policies or can build them
- You can pilot a single use case before going broad

## Reasons to wait or go slower

- Staff capacity is already stretched to the limit
- You serve vulnerable populations and haven't yet mapped what data those relationships involve
- Data infrastructure or policies are unclear
- Leadership is resistant or skeptical without engagement
- No budget exists for even basic tools or training

## LONG-TERM COSTS AND BENEFITS

Most AI tools charge monthly subscriptions: \$20 to \$300+ per user depending on the platform. Factor in staff time for setup and ongoing prompt refinement, periodic review of AI outputs, and potential costs for data storage or integrations. The payoff is real but not instant: most nonprofits report meaningful time savings within 3 to 6 months of consistent use, with the biggest gains in writing-heavy roles. Budget for the first year as a learning investment, not a solved problem.

## Employee buy-in: the piece most organizations skip

### 1 Name the fear

Give staff a real forum to voice concerns. Job security, accuracy, client privacy. None of these go away if you ignore them.

### 2 Involve before you decide

The people doing the work know where the time is lost. Include them in tool selection, not just rollout.

### 3 Train, don't announce

A memo isn't training. Budget real time for hands-on practice. Include people who are skeptical: they often become the best critical evaluators.

### 4 Keep the feedback loop open

Set a 60-day check-in. What's working? What isn't? Who feels left behind? Adjust before problems become resentment.

## IF YOUR STAFF ARE ALREADY USING AI ON THEIR OWN

Many organizations arrive at this guide because someone noticed staff are already using ChatGPT, Claude, or similar tools informally: on personal accounts, on personal devices, without any organizational policy in place. That's the more common starting point, and it changes what you need to do first.

The risk isn't that staff were curious and initiative. The risk is that client data, donor information, internal financials, or other sensitive material may have already passed through a third-party AI system you haven't reviewed or approved: with no record of what happened to it.

### Do these four things before anything else:

1. **Find out what's already in use.** Ask: without blame. You need an honest picture of which tools, how often, and for what tasks. You won't get it if staff feel they're being investigated.
2. **Assess what data went in.** If anyone entered client names, case details, donor records, immigration status, health information, or financial data into a free consumer AI tool, review that tool's privacy terms immediately. Most free tiers retain inputs for model training. That data is gone and you can't retrieve it, but you need to know whether a breach disclosure obligation exists.
3. **Issue an interim policy, even a short one.** One page is enough: what data may not enter any AI tool, what tools are currently approved (none, until you review them), and that staff should flag questions to a named person. This document stops the bleeding while you build something more complete.

4. **Then work through this guide in order.** Informal use doesn't mean you've failed: it means the organization has energy for AI adoption and you need a structure to channel it. That's a better starting position than resistance.

## What to do before you come back to this

If the readiness section gave you pause, pay attention to that. Adopting AI before the preconditions are in place creates more problems than it solves. Here's what to work on first.

### Precondition 01: Get your data house in order

Before you put organizational data into any AI tool, you need to know what data you have, where it lives, and who is allowed to access it. This doesn't require a formal data audit: it requires a honest inventory. Spend two hours with your team mapping out where client data, donor data, and staff data currently sit. That conversation alone will surface decisions you need to make before AI enters the picture.

### Precondition 02: Name an internal lead

AI adoption without an internal champion almost always stalls. This person doesn't need technical expertise: they need credibility with staff, curiosity, and time. Even two to three hours a week is enough in the early stages. If no one on your current team can take this on, that's a real constraint. Don't launch without it.

### Precondition 03: Write a one-page AI use policy

It doesn't have to cover everything. It needs to answer three questions: What data can staff put into AI tools? What types of decisions require human review? Who do staff contact if something goes wrong? One page. Approved by leadership before any tool goes live. The Fast Forward Policy Builder ([ffwd.org](http://ffwd.org)) can generate a starting draft in under 20 minutes.

### Precondition 04: Pick exactly one use case

The most common early mistake is trying to use AI everywhere at once. Pick the single task where your team loses the most time to repetitive writing. Grant reports. Volunteer onboarding emails. Social media captions. One thing. Run it for 30 days. Measure the time saved. Then decide whether to expand. Scope is your friend at the start.

**When to revisit:** Once you have a data inventory, an internal lead, a one-page policy, and one clear use case in mind: come back to the full guide. Those four things take most organizations two to four weeks to sort out. That's not a long delay. It's the difference between an implementation that works and one that creates problems you'll spend months cleaning up.

## When to keep AI out entirely

The planning categories section notes that AI shouldn't touch "direct client support or crisis response." That's right, but it understates the issue for some organizations. Certain populations and use cases require a harder line.

### Immigration services

- Intake screening or eligibility determinations
- Any document that will be submitted to USCIS or an immigration court
- Communications that could affect case status
- Anything involving legal advice or legal interpretation

### Mental health and crisis services

- Crisis line responses or risk assessment
- Therapy notes or clinical documentation
- Any communication with someone in acute distress
- Safety planning or discharge decisions

### Domestic violence and survivor services

- Intake or screening conversations
- Safety planning
- Any communication that could be discovered by an abuser
- Documentation that informs legal proceedings

### Child welfare and youth services

- Mandatory reporting decisions or documentation
- Case notes that inform placement or custody decisions
- Direct communication with minors without supervision
- Anything subject to FERPA or child welfare confidentiality rules

**The underlying principle:** AI is statistically reliable but not individually reliable. It performs well across a large population of inputs and fails unpredictably on edge cases. In any situation where a single error has serious

consequences for a specific person: legal status, physical safety, mental health, child welfare: the stakes are too high to accept statistical reliability. Keep humans in the decision seat.

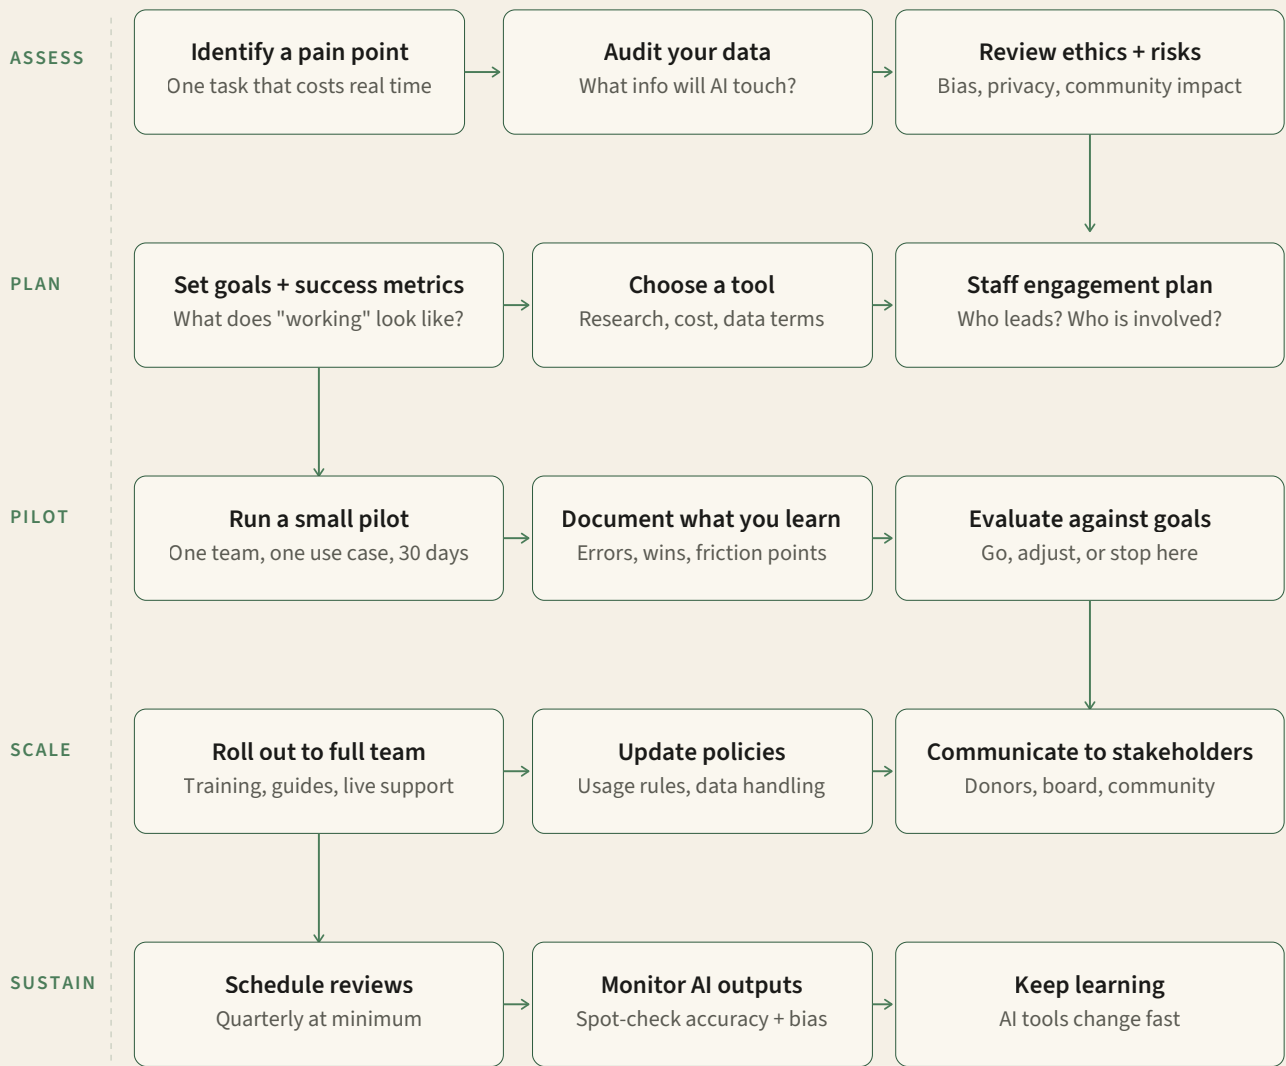
PART 2

## Getting started

The action sequence: from process to planning to tools to vendor scrutiny to staff readiness.  
These sections are meant to be worked through in order.

# Step-by-step process

Follow this sequence. Each phase builds on the last. Skipping ahead is the most common reason implementations stall.



# Planning categories

These aren't one-and-done. Revisit each area as your use of AI evolves.

## Start here: low-stakes use cases that build staff confidence

Before moving to anything complex or client-facing, start with tasks that are low-risk, easy to verify, and immediately useful. These build familiarity with how AI tools actually behave and give staff a chance to develop good habits around review and correction before the stakes are higher.

- **Meeting notes and action item summaries:** paste in a transcript or rough notes, ask for a clean summary. Easy to verify against what you know happened.
- **Donor thank-you letter drafts:** AI handles the first draft; staff review and personalize before anything goes out.
- **Social media post drafts:** low consequence if the first version isn't right, and a good way to see how much editing your team actually needs to do.
- **FAQ and volunteer onboarding documents:** AI can generate a solid first draft from existing materials; staff verify accuracy against what they know to be true.
- **Board report summaries:** condense program updates or data into readable summaries. The people reviewing them will catch errors quickly.
- **Adapting existing content for a new audience:** take a staff-facing document and ask AI to rewrite it for volunteers, or a technical report for a general donor audience. Low risk because the source material is already yours and already verified.

## CATEGORY 01 **Data and privacy**

---

What information will AI systems access, store, or generate. Who governs it?

- Inventory what client, donor, and staff data currently exists and where it lives  
Spreadsheets, CRMs, intake forms, email: all of it.
- Determine what data should never enter an AI tool  
Health records, immigration status, financial details, anything subject to HIPAA or other regulations.
- Review the data privacy terms of any tool before signing up  
Where is your data stored? Is it used to train the model? Can you delete it?
- Write or update your organization's AI data use policy  
Staff need clear rules about what they can and can't put into AI tools.
- Establish a repeatable process for preparing clean versions of documents before AI use  
Scrub or redact sensitive fields before input (client names, case numbers, identifying details) rather than feeding raw files. This will happen regularly. The process needs to be documented and consistently applied, not handled ad hoc each time.

Physical security considerations are covered in [Section C](#): screen visibility, shared workstations, and how to handle printed outputs.

## CATEGORY 02 **Ethics and equity**

---

Who benefits from this tool. Who might be harmed by it?

- Assess whether the tool could introduce or amplify bias  
AI trained on biased data produces biased outputs. Ask vendors directly about this.
- Consider how AI use might affect trust with your community  
Some populations have legitimate reasons to distrust automated systems.
- Establish a clear human review process for AI-generated content  
Especially anything that affects individual clients or goes out publicly.
- Disclose AI use to your community where appropriate  
Transparency isn't weakness: it's accountability.

## CATEGORY 03 **Staff and capacity**

---

Does your team have what they need to use AI well. Can they push back when it's wrong?

- Identify who will lead AI adoption internally  
This person doesn't need to be a tech expert: they need to be trusted and curious.
- Plan real training time, not just a handout  
Budget at least 2 to 4 hours of hands-on time per staff member at launch.
- Set clear expectations about human oversight  
AI writes the draft. A person reads it, edits it, and takes responsibility for it.
- Create a way for staff to flag problems  
If the AI is producing something wrong or harmful, staff need a clear path to report it.

## CATEGORY 04 **Operations and workflow**

---

How does AI fit into the work you already do. Where does it actually help?

- Map the tasks where AI will save the most time  
Writing, summarizing, answering repetitive questions, formatting data.
- Identify tasks AI should not touch  
Direct client support, crisis response, any decision with serious consequences.
- Document your workflows before and after AI integration  
This helps you train new staff and catch drift over time.
- Build quality checks into the workflow, not on top of it  
Review before publish, not as an afterthought.

## CATEGORY 05 **Financial planning**

---

What does this actually cost. How do you make the case to funders?

- Calculate total cost of ownership, not just subscription fees  
Staff time for training, setup, ongoing review, and eventual tool switching.
- Estimate the time savings in hours per week  
This becomes your ROI story for board and funders.
- Look for nonprofit discounts and grant funding  
Many AI providers offer discounted or free tiers for nonprofits. Some foundations now fund technology adoption.
- Plan for the tool to change  
AI pricing and features shift constantly. Don't build a budget that assumes nothing changes.

CATEGORY 06 **Governance and accountability**

---

Who is responsible when something goes wrong?

- Assign clear ownership over your AI policy  
Someone on staff: or the board: needs to be accountable for this.

---

- Set a regular review schedule  
At minimum, review your AI use quarterly in the first year.

---

- Define what would cause you to stop using a tool  
Know your line before you need it.

---

- Keep the board informed  
AI adoption is a governance matter, not just a staff decision.

## Tools to consider

This isn't a complete list and it will go out of date. It's a starting point. Most of these tools offer nonprofit discounts or free tiers: always ask before paying full price.

### ChatGPT (OpenAI)

FREE / \$25/SEAT/MO  
(BUSINESS, ANNUAL):  
NONPROFITS: UP TO 75%  
OFF

The most widely used general-purpose AI assistant. Strong at drafting, summarizing, editing, and answering questions. The free tier is genuinely useful for getting started. The Business plan (minimum 2 seats) keeps your data off OpenAI's training sets.

Best for: Grant writing drafts, donor communications, staff Q&A, internal documents.

### Claude (Anthropic)

FREE / \$25/SEAT/MO (TEAM  
STANDARD, ANNUAL: MIN. 5  
SEATS)

Strong on nuanced writing, long documents, and careful reasoning. Tends to handle sensitive topics more cautiously than other tools, which may matter for organizations working with vulnerable communities.

Best for: Policy drafts, ethics-sensitive writing, summarizing long reports.

### Gemini (Google)

INCLUDED IN GOOGLE  
WORKSPACE: BUSINESS  
STANDARD \$14/USER/MO  
(ANNUAL)

Since January 2025, Gemini is bundled into all Google Workspace Business and Enterprise plans: no longer a separate add-on. Integrates directly with Gmail, Docs, Sheets, and Meet. If your organization already pays for Workspace, you likely already have access.

Best for: Organizations on Google Workspace who want AI built into existing workflows.

### Copilot (Microsoft)

\$21/USER/MO (COPILOT  
BUSINESS, UP TO 300  
USERS, ANNUAL)

The Microsoft equivalent: built into Word, Excel, Outlook, and Teams. Most useful if your organization already runs on Microsoft 365. Nonprofit licensing discounts are often significant.

Best for: Organizations on Microsoft 365. Drafting in Word, summarizing in Outlook, analyzing data in Excel.

### Canva AI

FREE FOR NONPROFITS: UP  
TO 50 SEATS AT NO COST

Design tool with built-in AI features for generating images, writing social captions, and resizing content. The Canva for Nonprofits program gives eligible organizations full Canva Pro access: including all AI tools: completely free for teams of up to 50 people. Apply at [canva.com/canva-for-nonprofits](https://canva.com/canva-for-nonprofits).

Best for: Social media content, flyers, event materials, annual report visuals.

### Notion AI

INCLUDED IN NOTION  
BUSINESS: \$20/USER/MO  
(ANNUAL)

If your organization uses Notion for internal knowledge management or project tracking, AI is now bundled into the Business plan rather than sold as a separate add-on (the add-on was discontinued in May 2025). Lets you summarize, draft, and search across your existing content. Not a standalone tool.

Best for: Organizations already on Notion who want AI inside their knowledge base.

**On nonprofit pricing:** Microsoft, Google, Salesforce, and Canva all have established nonprofit programs. Canva's is particularly strong: full Pro access for up to 50 users at no cost. TechSoup is the standard clearinghouse for verifying eligibility and accessing Microsoft and Salesforce offers. Before budgeting full price for any major tool, check [techsoup.org](https://techsoup.org) and the vendor's nonprofit page directly.

### FREE RESOURCE

**Fast Forward's Nonprofit AI Policy Builder**, a free tool that generates a custom AI policy for your organization based on your answers to a short set of questions. A good starting point if you need a policy document and don't know where to begin. Find it at [ffwd.org](https://ffwd.org).

## BUILT SPECIFICALLY FOR NONPROFITS

The tools above are general-purpose. These two are purpose-built for nonprofit grant work, with security credentials and nonprofit-specific pricing to match. If grant writing is where your team spends the most time, start here instead.

### Grantable

FREE TIER / \$75/MO  
(UNDER \$500K BUDGET)

Grant writing platform built exclusively for nonprofits, institutions, and grant agencies. Learns your organization's voice from past proposals and generates drafts that match your tone. Your content is never used to train AI models, and the platform holds SOC 2 Type 2 certification with third-party security audits.

Best for: Organizations where grant writing is a significant staff burden. Especially useful for small teams applying to multiple funders with similar narrative requirements.

[grantable.co](https://grantable.co)

### Instrumentl

FROM \$299/MO: 14-DAY  
FREE TRIAL

Full grant lifecycle platform covering discovery, writing, tracking, and reporting. The AI drafting tool draws on a database of 400,000+ funder profiles, which means it can tailor proposals to specific funder priorities in ways a general-purpose AI can't. Trusted by 5,500+ nonprofits. Secure storage and version history built in.

Best for: Development teams managing multiple active grants. The higher price makes most sense if you're already spending significant staff hours on grant research and tracking.

[instrumentl.com](https://instrumentl.com)

## GO DEEPER: AI AND GRANT WRITING

For workflow guidance, risks, and how to use AI in grant writing without the pitfalls, see [Section F: AI and grant writing](#).

# Questions to ask every vendor

Most AI tools don't volunteer this information. You have to ask. These five questions are non-negotiable before you put any organizational or client data into a new tool.

- **Does this product train on my data?**

Some tools use your inputs to improve their models. If you're putting donor information, client intake notes, or internal strategy documents into a tool, you need to know whether that content is being used to train a system that other users benefit from. Ask for the specific policy in writing, not a general privacy overview.

- **Where is my data stored, and who can access it?**

Data residency matters for organizations subject to state or federal regulations. "Stored in the cloud" is not an answer. Ask for the specific region, the access controls, and whether any third-party subprocessors have access to your data.

- **Can we delete our data, and how?**

If you end your subscription or want to remove specific content, what's the process? Is deletion immediate or does data persist in backups? Organizations working with sensitive populations need a clear, documented deletion path before they commit to a platform.

- **What is your security certification: SOC 2, HIPAA, or equivalent?**

If your organization handles health information, financial data, or serves regulated populations, the tool needs to meet minimum security standards. SOC 2 Type II is the standard benchmark for most SaaS tools. HIPAA compliance is required if health data is involved. No certification should mean no deal for those use cases.

- **Do you offer a nonprofit or mission-driven discount?**

Many AI vendors offer substantial discounts that aren't advertised on the pricing page. The worst they can say is no. Ask before you budget, not after you've already committed.

**Red flag:** If a vendor can't answer these questions directly, or points you to a general terms of service page instead of giving a clear yes or no, treat that as meaningful information. Transparent vendors have clear answers ready. Evasion usually means the answer is one you wouldn't like.

# Skills, hiring, and training

You don't need a data scientist to use AI responsibly. But you do need specific skills in place before you launch. You also need a clear-eyed view of whether you're building them internally, hiring them, or bringing them in on contract.

## REQUIRED Skills you need before you go live

Every organization adopting AI needs someone who can do these things. Not an expert: someone competent and accountable.

- Prompt writing: knowing how to give an AI tool clear, specific instructions**  
This is not a technical skill. It's a communication skill. Someone who writes well and thinks clearly can learn the basics in a few hours. The difference between a vague prompt and a specific one is usually the difference between a useless output and a useful one.
- Output evaluation: reading AI content critically and catching errors**  
AI sounds confident even when it's wrong. The person reviewing outputs needs to know your organization's programs, data, and voice well enough to spot when something is off. This is usually an existing staff member, not a hire.
- Basic data hygiene: knowing what should and shouldn't go into an AI tool**  
This is policy enforcement as much as skill. Someone needs to understand the data privacy policy well enough to apply it in practice and answer staff questions about edge cases.
- Vendor evaluation: reading a data privacy policy and asking the right questions**  
Covered in the vendor questions section of this guide. Not a technical skill: it requires knowing what to look for and being willing to push for clear answers.

## USEFUL BUT NOT REQUIRED **Skills that help as you scale**

These matter more once you're past a basic pilot and starting to integrate AI into core workflows.

### Understanding of how AI models work at a basic level

Not programming: just a working mental model of what these systems can and can't do, why they hallucinate, and why bias enters. Enough to explain it to a skeptical board member or a worried community partner.

### Workflow documentation

Writing down how AI fits into existing processes so new staff can be trained and so you can catch drift over time. Usually a natural fit for someone already doing operations or communications work.

### Data literacy: reading a spreadsheet, interpreting a chart, asking good questions about data

If you're using AI for donor analysis or program reporting, someone needs to be able to evaluate whether the outputs make sense. This isn't about building models: it's about not accepting outputs uncritically.

### Change management

Getting skeptical staff on board, handling the fear and frustration that comes with any new system, and keeping morale intact through the transition. Often the ED or a trusted program manager rather than a technical person.

## **Build, hire, or contract: how to decide**

**Build internally**

**BEST FOR MOST ORGS**

Most of the skills above can be developed in existing staff with 4 to 8 hours of focused training. Identify one person with strong writing instincts, good judgment, and institutional knowledge. Give them time to learn and make them your internal lead. This works for the majority of nonprofits at the pilot stage.

Good fit when: You have a capable generalist on staff and a single, low-stakes use case to start with.

**Hire a specialist**

**JUSTIFIED AT SCALE**

A dedicated technology or AI role makes sense when AI is embedded in multiple workflows, when your data infrastructure needs attention, or when the volume of oversight work exceeds what any existing staff member can absorb alongside their current job. This is a later-stage decision, not a precondition for starting.

Good fit when: AI has become operationally central and the internal lead is stretched beyond capacity.

**Bring in a consultant**

**BEST FOR SPECIFIC GAPS**

A consultant or fractional AI specialist is the right call for one-time tasks: setting up your AI governance framework, auditing your data privacy practices, training your full staff, or evaluating tools before you commit. Not a substitute for internal capacity. A way to fill gaps without a permanent hire.

Good fit when: You need deep expertise for a defined project and don't need it ongoing.

**Free and low-cost training resources**

**Google for Nonprofits AI training**

Free access to Google's AI literacy curriculum through the Google for Nonprofits program. Covers generative AI basics, responsible use, and practical applications. Start at [google.org/nonprofits](https://google.org/nonprofits).

**Microsoft Nonprofit AI learning paths**

Free courses on AI fundamentals, responsible AI, and Microsoft Copilot through Microsoft Learn. Eligible organizations can access these via the Microsoft Nonprofit program at [microsoft.com/nonprofits](https://microsoft.com/nonprofits).

**TechSoup learning library**

Webinars, guides, and training resources on technology adoption for nonprofits, including AI. Many are free or low-cost for verified nonprofits. [techsoup.org/learning](https://techsoup.org/learning).

**Fast Forward AI resources**

Practical guides, case studies, and frameworks written specifically for nonprofit leaders. Also publishes the annual AI for Humanity report. [ffwd.org](https://ffwd.org).

**LinkedIn Learning (via Microsoft nonprofit discount)**

**NTEN (Nonprofit Technology Enterprise Network)**

Access to thousands of courses including AI fundamentals, prompt engineering, and responsible AI. Organizations on Microsoft 365 nonprofit licensing often have access included.

Community of nonprofit technology practitioners. Offers training, peer learning, and an annual conference focused on technology in the sector. [nten.org](https://www.nonprofittechnology.org).

### PART 3

## Doing it responsibly

The ongoing accountability layer. Once AI is live, these are the questions that don't go away: oversight, measurement, communication, client protection, contracts, and policy.

# What human oversight actually looks like

"Maintain human oversight" appears throughout this guide. Here's what that means in practice, for three common nonprofit use cases.

## Use case: AI-assisted grant writing

- 1 Staff member drafts a prompt with specific program data, outcomes, and funder priorities.
- 2 AI generates a first draft. Staff member reads the full draft for accuracy: every statistic, every claim about the program, every description of community need.
- 3 Staff member edits for voice, accuracy, and alignment with the organization's actual work. AI drafts often generalize; the editor's job is to make it specific and true.
- 4 Program director or ED reviews the final version before submission. They are accountable for what gets submitted: not the AI tool.
- 5 Any grant that misrepresents program outcomes is a compliance issue regardless of how the draft was produced.

## Use case: Donor email or newsletter

- 1 Communications staff drafts a prompt with key messages, audience, tone, and any specific asks.
- 2 AI produces a draft. Staff reads for accuracy, tone, and whether it actually sounds like the organization.
- 3 Check specifically for hallucinated details: dates, names of programs, funding amounts, impact numbers. AI fills gaps with plausible-sounding content that may be wrong. ([→ glossary](#))
- 4 A second staff member or supervisor reviews before send, the same as any external communication.
- 5 The communications policy governs this content. AI is a drafting tool, not an approval process.

### Use case: Summarizing reports or meeting notes

---

- 1 Staff member uploads or pastes a document and prompts the AI to summarize key points.
- 2 Staff member reads the summary against the original document. Summaries can omit important nuance, mischaracterize positions, or drop critical details.
- 3 If the summary will be shared externally or used in decision-making, the person sharing it is responsible for its accuracy: not the tool that generated it.
- 4 Internal use with low stakes (a quick recap of a meeting) carries lower risk than a summary that will inform a board decision or funder report.

**The test is simple: could a real person defend this output if asked? If the answer is yes: because they read it, checked it, and stand behind it: oversight is working. If the answer is "the AI said it," oversight has failed.**

# How to know if it's working

AI adoption succeeds or fails quietly. Without clear metrics from the start, you'll reach month six without being able to answer whether it was worth it. Decide what you're measuring before you launch: not after.

## Time metrics

Hours per week spent on grant writing, before and after

Time to produce a first draft of a donor email or newsletter

Time spent on meeting summaries and internal documentation

Hours reclaimed per staff member per week: cumulative across the team

## Capacity metrics

Number of grant applications submitted per quarter, before and after

Number of donor communications sent per month

Volume of translated or adapted materials produced

Tasks completed that previously went undone due to bandwidth

## Staff metrics

Staff confidence with the tool at 30, 60, and 90 days (simple 1 to 5 survey)

Number of staff actively using the tool vs. number trained

Problems or concerns flagged through the feedback channel

Whether skeptical staff have shifted, and if not, why not

## Quality metrics

Grant award rate before and after (give this 6 to 12 months to be meaningful)

Donor response rates to AI-assisted communications

Errors or corrections caught during human review: track the pattern

Whether outputs consistently reflect your organization's voice and values

**The 90-day check-in question:** If you could only ask your team one thing at 90 days, ask this: "Is AI saving you time or creating work?" If the honest answer is "creating work," something in the implementation needs to change before you expand. That's not failure. That's useful data.

## What to tell your board, funders, and community

AI adoption is increasingly visible: to funders who are starting to ask about it, to board members who are reading about it, and to communities who have legitimate reasons to want to know. Getting ahead of these conversations is easier than responding to them under pressure.

### AUDIENCE 01 **Your board**

What they need to know, and when.

- Brief the board before launch, not after  
Frame it as a governance decision: which tools, which use cases, what oversight is in place, what's off-limits. Give them a chance to ask questions before it's operational.
- Report on it at least quarterly in the first year  
What's working, what isn't, any problems that surfaced. Boards that feel informed are far more likely to support the work.
- Be honest about what you don't know yet  
"We're piloting this and will have more data in 90 days" is a better answer than overpromising outcomes you can't yet demonstrate.

## AUDIENCE 02 Funders

---

An evolving conversation, and one worth starting proactively.

- Know your funder's AI policy before you submit

A small number of funders have explicit policies about AI use in grant applications. Most haven't addressed it yet. If you're unsure, ask your program officer directly: it builds trust rather than undermining it.

- Disclose AI involvement in proposals when it's significant

If AI drafted substantial portions of a proposal, say so briefly. Something like: "Drafts were developed with AI assistance and reviewed and edited by staff." This is accurate, professional, and increasingly expected.

- Lead with what stayed human

The program design, the community relationships, the theory of change: that's your staff's work. AI handled the formatting and language efficiency. Make that distinction clear.

- Use AI efficiency as a capacity story

If AI allows your two-person development team to apply to three times as many grants, that's a legitimate organizational strength to share with funders who care about sustainability.

## AUDIENCE 03 Your community

---

The people you serve have the most at stake. Treat that accordingly.

- Be transparent about where AI is involved in communications

If your newsletter or donor emails are AI-assisted, that doesn't need to be front-page news: but it shouldn't be hidden either. A brief note in your communications policy or website is enough for most organizations.

- Never use AI in direct client-facing interactions without disclosure

If a chatbot or AI tool is responding to community members on your behalf, they need to know they're not talking to a person. This is a basic trust and consent issue.

- Listen for community skepticism and take it seriously

Some communities: particularly those that have experienced algorithmic discrimination or surveillance: will have strong feelings about AI. Those concerns are legitimate. Address them directly rather than dismissing them.

## When AI faces your clients directly

Using AI internally to draft documents, summarize reports, or support staff is a fundamentally different decision than deploying AI to interact with the people you serve. The error tolerance is lower. The oversight requirements are higher.

### COMMON USE CASES **What nonprofits are deploying and where**

- Chatbots for intake and initial screening**  
Automated tools that ask initial eligibility questions, collect contact information, or triage service requests before connecting clients to staff. These reduce intake burden for small teams but introduce risks around accuracy, accessibility, and what happens when the chatbot is wrong.
- Translation tools for multilingual client communications**  
AI translation has improved significantly and can help organizations reach non-English-speaking clients more quickly and consistently than before. The risk is over-reliance on machine translation for sensitive topics. Legal status, health information, safety planning. Nuance and accuracy are non-negotiable in those contexts.
- Automated scheduling and reminders**  
Lower-stakes and generally appropriate for most populations. The main considerations are accessibility (not all clients have smartphones or reliable internet) and whether the communication style fits your population.
- AI-assisted case documentation**  
Staff using AI to summarize client notes or generate documentation templates. This is internal use even though it involves client information: but it carries the same data privacy obligations as any other client data handling.

## POPULATIONS THAT REQUIRE EXTRA CARE **When the stakes for error are higher**

For some populations, an AI system that makes an error: misclassifies eligibility, mistranslates critical information, or fails to recognize a safety concern: can cause direct harm. These are not reasons to avoid AI, but reasons to require a higher level of oversight and testing before deployment.

### Clients in crisis or safety situations

Domestic violence survivors, individuals in mental health crisis, people experiencing homelessness during extreme weather. These situations require human judgment. An AI chatbot that misreads urgency or follows a script when the situation has escalated beyond the script is a liability, not an asset.

### Clients with limited English proficiency

Machine translation can fail in ways that aren't obvious to English-speaking staff reviewing the output. Have translated materials reviewed by a fluent speaker before using them, especially for anything involving rights, safety, or eligibility.

### Clients with disabilities

Screen readers, cognitive accessibility, low literacy: AI-generated content isn't automatically accessible. Test any client-facing tool with actual users from your population before wide deployment.

### Clients with past trauma involving automated systems

Some populations: particularly communities that have experienced discriminatory screening in housing, employment, or benefits: have legitimate reasons to distrust automated systems. Transparency about what is AI-generated and what involves a human can help. So can giving clients a way to opt out.

BEFORE YOU DEPLOY **What responsible client-facing AI looks like**

---

Test with real users from your population before launch

Not staff testing on behalf of clients. Actual users, with actual needs, interacting with the actual tool. The failure modes that matter are the ones your clients encounter.

---

Build a clear human escalation path

Every client-facing AI system needs a well-marked, easy-to-access way to reach a human. If a client can't find it, it doesn't count.

---

Disclose that AI is involved

Clients have a right to know when they're interacting with an automated system. This is both an ethical baseline and, depending on your state's evolving privacy law, potentially a legal one.

---

Define and monitor failure conditions

What does the AI do when it doesn't understand a question? What happens if it produces a wrong answer about eligibility or services? Who reviews that interaction? Know the answers before you launch.

# Procurement and vendor contracts

Many nonprofits: especially those with government contracts or specific funder requirements: go through a formal procurement or RFP process for technology adoption. AI tools require some additions to standard technology procurement language.

## RFP LANGUAGE TO ADD **What to require AI vendors to address in proposals**

### Data use and training disclosure

Require vendors to state explicitly whether customer data is used to train their models, under what conditions, and whether customers can opt out. "We take privacy seriously" is not an answer to this question.

### Data retention and deletion policies

How long does the vendor retain your data? What happens to it when you cancel? Can you request deletion and within what timeframe? These questions need specific, written answers.

### Subprocessors and third-party access

Who else has access to your data? Many AI tools rely on cloud infrastructure or third-party model providers. You need to know the full chain.

### Compliance certifications

Require vendors to identify current certifications: SOC 2 Type II, HIPAA Business Associate Agreement (if applicable), FedRAMP, or sector-specific compliance frameworks your funder requires. Ask for documentation, not just claims.

### Incident notification procedures

If there is a data breach, how quickly will you be notified? What is the vendor's process? This matters for your own reporting obligations.

### Model change notification

AI tools update their underlying models regularly. Require vendors to notify you of material changes that could affect output quality, safety filtering, or data handling: not just feature changes.

## CONTRACT CLAUSES TO WATCH FOR **Terms that should give you pause**

### Broad license grants to your data

Some AI tool contracts include language granting the vendor a broad license to use your content and data for unspecified purposes. Read the data use terms carefully, not just the privacy policy.

### Unilateral terms changes

Clauses allowing the vendor to change terms with minimal notice, or that define continued use as acceptance of new terms. Push for advance notice requirements and the right to exit on material changes.

### Limitation of liability provisions

AI tools sometimes disclaim liability for outputs entirely. Understand what you're agreeing to when something goes wrong. Consider whether your organization can absorb that risk.

### Automatic renewal with long notice windows

Subscription contracts that auto-renew annually and require 60-90 day cancellation notice. Set a calendar reminder the moment you sign.

**Before you sign:** Many AI vendors, including well-known consumer tools, publish extensive terms of service designed for consumer use that do not address organizational compliance needs. If a vendor can't produce a Business Associate Agreement, a Data Processing Agreement, or a straight answer about whether your data trains their model, that's information. Small nonprofits rarely have bargaining power for custom terms, but you always have the option to choose a different vendor.

# AI use policy template

Every organization using AI needs a written policy. What follows is a template you can adapt. Read it in full, edit it to reflect your actual situation, and have your executive director and board chair review it before it goes into effect. A policy no one has read is not a policy.

TEMPLATE: EDIT BEFORE USE

## [Organization Name]

### Artificial Intelligence Use Policy

Effective date: [DATE] | Next review: [DATE, recommend annually] | Policy owner: [NAME/ROLE]

#### 1. PURPOSE

[Organization Name] uses artificial intelligence tools to support staff in their work. This policy establishes how AI tools may be used, what information may and may not be entered into these tools, and who is responsible for oversight. The goal is to capture the efficiency benefits of AI while protecting the people we serve, our donors, our staff, and the organization's integrity.

#### 2. SCOPE

This policy applies to all staff, volunteers, and contractors who use AI tools: including but not limited to ChatGPT, Claude, Microsoft Copilot, Google Gemini, and any other generative AI product: for work-related purposes, whether on organizational devices or personal devices used for work.

#### 3. APPROVED TOOLS

The following AI tools have been reviewed and approved for organizational use:

[List approved tools here with any conditions, e.g.: "Claude Pro (Anthropic): approved for drafting and summarization. Do not enter client data. Team account managed by [NAME]." If no tools have been

formally approved yet, write: "No tools are currently approved for organizational use. Staff should contact [NAME] before using any AI tool for work-related purposes."

Use of AI tools not on this list for work-related purposes requires prior approval from [NAME/ROLE]. Staff who are uncertain whether a tool requires approval should ask before using it.

---

#### 4. WHAT MAY NOT BE ENTERED INTO AI TOOLS

The following categories of information may not be entered into any AI tool under any circumstances, unless that tool has been specifically reviewed and approved for that data type:

- Client names, identifying information, case details, or service records
- Health, mental health, or disability information
- Immigration status or documentation
- Donor names, giving history, or financial information
- Staff personnel records, salaries, or performance information
- Financial records, account numbers, or audit materials
- Any information subject to a confidentiality agreement, grant restriction, or legal obligation

When in doubt about whether information is safe to enter, do not enter it. Ask [NAME/ROLE] first.

---

#### 5. HUMAN REVIEW REQUIREMENT

All AI-generated content that will be used externally: in grant proposals, donor communications, public materials, client-facing documents, or official correspondence: must be reviewed and approved by a staff member before use. The reviewing staff member is responsible for the accuracy and appropriateness of the final content. "AI generated it" is not a defense for an error.

AI-generated content used for internal purposes only (meeting summaries, internal drafts, research notes) does not require formal sign-off but should still be read and verified by the person using it.

---

#### 6. DISCLOSURE

[Organization Name] will disclose its use of AI when asked by clients, donors, funders, or partners. Staff should be prepared to honestly answer questions about whether AI tools were used in a document or communication. [Add any funder-specific disclosure requirements here, e.g.: "Grant proposals submitted to [Funder Name] must include a disclosure statement if AI was used in drafting."]

## 7. REPORTING CONCERNS

Staff who believe an AI tool produced harmful, inaccurate, or biased content, or who are uncertain whether a particular use of AI is appropriate under this policy, should report to [NAME/ROLE]. No one will be penalized for raising a concern in good faith.

## 8. POLICY OWNER AND REVIEW

This policy is owned by [NAME/ROLE] and will be reviewed [annually / every six months] or whenever a material change in the organization's AI use warrants an update. The board of directors will be informed of material changes to AI policy and practice.

Executive Director signature

Board Chair signature

\_\_\_\_\_  
Name | Date

\_\_\_\_\_  
Name | Date

**After you adopt this policy:** Give every staff member a copy. Review it at onboarding for new hires. Revisit it when you add a new AI tool, when a vendor changes their terms, or when something goes wrong. A policy that lives in a folder and never gets looked at is not protecting anyone.

PART 4

# When you need it

Pull these up when you need them. Not linear reading: reference material for specific questions: terms, common mistakes, security, local vs. cloud, and Minnesota-specific resources.

# Terms you'll encounter

You don't need to be a technologist to use AI tools responsibly. But you do need to know what people mean when they use these words.

## Artificial intelligence (AI)

Software that performs tasks typically requiring human judgment: writing, summarizing, answering questions, recognizing images. It doesn't "think." It produces outputs based on patterns learned from large amounts of data.

## Large language model (LLM)

The technology behind tools like ChatGPT and Claude. Trained on massive amounts of text to predict and generate human-like language. It has no understanding of meaning: it works on statistical patterns.

## Prompt

The instruction or question you give an AI tool. "Write a thank-you letter to a major donor" is a prompt. The quality of your prompt has a major effect on the quality of the output.

## Training data

The content an AI model learned from. Most general-purpose tools were trained on large portions of the internet, books, and other text. This is where bias enters: models reflect the biases present in their training data.

## Hallucination

When an AI generates something that sounds confident and factual but is wrong or made up. Names, statistics, citations, and dates are especially prone to this. Always verify before using AI-generated facts publicly.

## AI-generated content

Any text, image, or other output produced by an AI tool. The human who prompted the tool is still responsible for reviewing, editing, and approving it before it goes anywhere.

## Model

The underlying AI system a tool is built on. ChatGPT runs on GPT-4. Claude runs on Anthropic's Claude model. Different models have different strengths, limitations, and policies.

## Bias

When AI outputs reflect unfair patterns: consistently producing worse results for certain groups, reinforcing stereotypes, or making systematically different errors based on race, gender, language, or other characteristics.

### **Human oversight**

A designated person reviewing AI outputs before they're used, published, or acted on. Not a rubber stamp: an actual check for accuracy, appropriateness, and potential harm.

### **Data privacy policy**

Your organization's written rules about what information can and can't be entered into AI tools, how AI outputs should be handled, and who is responsible for enforcement. Every organization using AI needs one.

## What goes wrong in the first year

These aren't hypothetical risks. They're the patterns that show up repeatedly when nonprofits adopt AI without enough structure. None of them are catastrophic on their own. Several of them together can erode staff trust and community confidence in ways that are hard to walk back.

### MISTAKE 01 **Publishing AI outputs without review**

The most common and most damaging error.

- AI-generated content goes out with wrong statistics, wrong names, or wrong dates  
AI fills gaps with plausible-sounding content. A donor letter that references the wrong program or the wrong year is a trust problem, not just a typo.
- Grant proposals contain inaccuracies about the organization's own work  
This has compliance implications. A funded proposal is a contract. Misrepresentation: even unintentional: creates real liability.
- The fix is simple but requires discipline  
Every AI output gets a human read before it leaves the organization. Build this into the workflow, not as an afterthought. If there isn't time to review it, there isn't time to send it.

## MISTAKE 02 **Training that's really just an announcement**

---

A memo, a link, and a 30-minute demo is not training.

- Staff who aren't confident with the tool either avoid it or misuse it  
Both outcomes waste the investment. Avoidance means the tool doesn't get used. Misuse means outputs don't get reviewed properly.
- Skeptical staff get left behind  
These are often the staff with the most institutional knowledge and the most to contribute to responsible use. Losing their engagement is a real loss.
- Budget two to four hours of hands-on practice per staff member  
Not a demonstration: practice. They should leave knowing how to write a prompt, what to do when the output is wrong, and who to call with questions.

## MISTAKE 03 **Using AI in client-facing work without disclosure**

---

Particularly damaging for organizations serving communities with low institutional trust.

- Community members find out after the fact that they were interacting with AI  
In mental health, social services, and advocacy contexts, this discovery can feel like a betrayal. Functionally, it is one.
- AI chatbots deployed on websites without clear labeling  
If it's not obvious that a responder is automated, it needs to be labeled. This is both an ethical obligation and, in some contexts, a legal one.
- The standard is simple: if AI is involved in a client interaction, the client should know  
This doesn't require a lengthy disclaimer. It requires honesty at the point of contact.

## MISTAKE 04 **Choosing tools based on price alone**

---

The cheapest option is rarely the right one for organizations handling sensitive data.

- Free tools often have the weakest data privacy terms**  
Your data may be used to train the model. Storage and deletion policies may be vague or nonexistent. Read the terms before you sign up, regardless of cost.
- Consumer tools aren't the same as enterprise tools**  
The free version of a general-purpose AI tool is designed for individual use, not for organizations managing client data. Enterprise and nonprofit tiers typically include stronger data controls.
- Price should be one factor among several**  
Security certification, data policy, nonprofit-specific features, and vendor support matter more for organizational use than they do for personal use.

## MISTAKE 05 **No plan for when the tool changes**

---

AI tools change constantly. Pricing, features, terms of service, and underlying models all shift: sometimes without much notice.

- Organizations build workflows around a tool and then the tool changes significantly**  
A price increase, a feature removal, or a change in data policy can disrupt operations that have come to depend on the tool.
- Staff build skills on a specific interface that gets redesigned**  
This isn't a reason to avoid AI: it's a reason to document your prompts, your workflows, and your policies so they're portable.
- Build flexibility into your AI governance from the start**  
Review tools annually. Keep a short list of alternatives. Don't let a single vendor become so embedded that switching becomes impossible.

## Physical and data security

AI tools introduce new security considerations on top of the ones your organization already manages. Most of them aren't exotic. They're the basics: device hygiene, access controls, offboarding: applied with AI tools in mind. The organizations that get into trouble are usually not the ones that were hacked. They're the ones that left a door open through carelessness.

## PHYSICAL SECURITY **What happens in the room**

Screen visibility and physical access are the easiest vulnerabilities to close. They're also the most frequently ignored.

- Screen locking: every device should lock automatically after a short period of inactivity**  
Two minutes is reasonable for a shared or open office environment. Staff should also lock manually when stepping away. This applies to any device where AI tools are accessed, including personal phones used for work.
- Screen visibility in open offices and public spaces**  
If a staff member is using an AI tool to draft a grant proposal or review donor data, that screen is visible to anyone nearby. Positioning desks away from foot traffic, using privacy screens, and establishing norms around sensitive work in public spaces are all worth addressing explicitly.
- Shared workstations and shared logins**  
If multiple staff members use the same computer, every AI tool session should be logged in and out individually. Shared logins make it impossible to track who entered what data and create accountability gaps that matter when something goes wrong.
- Printed AI outputs**  
Anything printed from an AI tool: a grant draft, a donor list, a program summary: is physical data. Apply the same handling rules you would to any sensitive document. Shred rather than trash. Don't leave printed materials in common areas.
- Work conversations in public spaces**  
Staff discussing client cases, donor information, or program strategy while using AI tools in coffee shops, co-working spaces, or transit creates real exposure. Establish clear norms about where sensitive AI-assisted work should and shouldn't happen.

## DEVICE AND ACCESS SECURITY **Who has access to what, and how it's controlled**

The most common security failures in nonprofits aren't technical attacks. They're access control problems.

### Use organizational accounts, not personal ones, for all AI tools

When staff sign up for AI tools using personal email addresses, the organization loses visibility and control. If that staff member leaves, the account, along with everything in it, goes with them. All AI tool accounts should be created under organizational email addresses and managed centrally.

### Strong, unique passwords and a password manager

Reusing passwords across tools is one of the most common ways organizations get compromised. A password manager like Bitwarden (free for nonprofits) makes strong unique passwords manageable without requiring staff to memorize them.

### Two-factor authentication on every AI tool account

Most enterprise and nonprofit-tier AI tools support two-factor authentication. It should be enabled and required. This alone stops the majority of unauthorized account access attempts.

### Personal devices used for work: establish a clear policy

Staff checking AI tools on personal phones or home computers creates a security gap you can't fully close. At minimum, require that work-related AI use happens on devices with screen locks, current software updates, and no shared access with family members.

### Network security: public wifi and VPNs

Using AI tools on unsecured public wifi exposes credentials and data in transit. Staff working remotely should use a VPN, especially when accessing tools that contain or process sensitive organizational data. Many VPN services offer nonprofit pricing.

## OFFBOARDING **What happens when a staff member leaves**

Offboarding is where access control failures most often surface. It requires a specific checklist, not good intentions.

- Revoke access to all AI tools on the day of departure

Not the end of the week. The day of. This includes tools the staff member set up independently on organizational accounts. Maintain a running inventory of every tool and every account so nothing gets missed.

- Change shared passwords immediately

If any AI tool accounts or organizational credentials were shared: even informally: change them the day someone leaves. This applies to departures of all kinds, including friendly ones.

- Review what data the departing staff member had access to

If the person handled donor data, client information, or sensitive program content through AI tools, document what they accessed and confirm it was governed by your data policy. This matters for compliance and for your own records.

- Retrieve or remotely wipe organizational data from personal devices

If a staff member used personal devices for work, have a policy in place: agreed to at hiring: that allows you to remove organizational data from those devices upon departure. Without this agreement in place in advance, enforcement is difficult.

## INCIDENT RESPONSE **What to do when something goes wrong**

Every organization using AI tools needs a short, clear plan for the moment something goes wrong. The plan doesn't need to be long. It needs to exist before you need it.

### Define what counts as a security incident for your organization

Unauthorized access to an AI tool account. A staff member entering client health data into a tool that trains on user inputs. A data breach at a vendor. A printed output left in a public space. Define these scenarios in advance so staff know when to escalate.

### Name a single point of contact for security incidents

Staff need to know exactly who to call and how quickly. Not "email the ED": a named person, a phone number, and a clear expectation of response time. For most small nonprofits this is the ED or operations director.

### Know your notification obligations

Depending on your state and the type of data involved, a security breach may trigger legal notification requirements: to affected individuals, to state regulators, or both. Know these obligations before an incident, not after. Your legal counsel or state nonprofit association can help clarify what applies.

### Notify your vendors

If you suspect a breach involves a specific AI tool or platform, notify the vendor immediately. Most enterprise and nonprofit-tier tools have security response teams and documented incident reporting processes. Use them.

### Document everything

Date, time, what happened, who was involved, what data was affected, what actions were taken. This documentation matters for compliance, for your board, and for any legal or regulatory process that follows.

**A note on cyber insurance:** If your organization doesn't have cyber liability insurance, AI adoption is a reasonable prompt to look into it. Premiums for small nonprofits are often lower than expected, and coverage typically includes breach notification costs, legal fees, and some recovery expenses. Your existing insurance broker can advise on whether your current policy covers AI-related incidents or whether a separate rider makes sense.

## Local models vs. cloud tools

---

Most AI tools run on servers owned by a technology company. Your inputs travel over the internet, get processed on their infrastructure, and return as outputs. There's an alternative: running an AI model directly on hardware your organization controls, where data never leaves your building. For most nonprofits, cloud tools are the right starting point. But the tradeoff matters if you handle sensitive data.

## Cloud-based tools

HOW MOST NONPROFITS  
START

Your prompts and any data you include in them are sent to a third-party server, processed, and returned. The AI model lives on the vendor's infrastructure. You access it through a browser or app.

### ADVANTAGES

- + No hardware investment or IT infrastructure required
- + Access to the most capable, regularly updated models
- + Works on any device with a browser: no installation
- + Predictable subscription costs, nonprofit discounts available

### TRADEOFFS

- Your data leaves your organization's control during processing
- Subject to vendor's data retention, storage, and training policies
- Requires active internet connection
- Pricing and terms can change without much notice

## Local (on-device) models

FOR HIGHER-SENSITIVITY  
USE CASES

The AI model runs entirely on hardware your organization owns. No data is sent to an external server. Processing happens on a local computer or server, and outputs never leave your building unless you choose to share them.

### ADVANTAGES

- + Data never leaves your organization: complete privacy control
- + No vendor data policies to track or monitor for changes
- + Works offline: useful in low-connectivity environments
- + No ongoing subscription costs once set up

### TRADEOFFS

- Requires capable hardware: most local models need a modern computer with significant RAM or a dedicated GPU
- Local models are less capable than the best cloud models: outputs may require more editing
- Setup and maintenance requires technical comfort, not plug and play
- You are responsible for updates, security patches, and ongoing maintenance

## HARDWARE REALITY **What running a local model actually requires**

Local AI is not just downloading an app. The hardware requirements are real and worth understanding before you commit.

### ❑ RAM is the primary constraint

Most capable local models require 8 to 16GB of RAM at minimum. A computer with 4GB of RAM (common in older nonprofit hardware) will not run them useably. 16GB is a comfortable floor for smooth performance on smaller models. 32GB opens up more capable options.

### ❑ A dedicated GPU accelerates performance significantly

Running AI models on a CPU alone is slow. A dedicated graphics card (GPU): particularly NVIDIA cards with CUDA support: speeds up processing dramatically. This is the difference between a 30-second response and a 3-second one. Not strictly required, but noticeable.

### ❑ Storage for the model files

Local AI models are large files: anywhere from 4GB to 40GB+ depending on the model. You'll need available disk space and enough remaining for normal operations. An SSD rather than a spinning hard drive improves load times significantly.

### ❑ The hardware investment may exceed a year of cloud subscriptions

A computer capable of running local models well costs \$800 to \$2,000+. For most nonprofits, this only makes financial sense if the privacy requirements are strong enough to justify it, or if you're already planning a hardware refresh.

## WHEN LOCAL MAKES SENSE **The use cases that justify the complexity**

Local models are not the right starting point for most nonprofits. They become worth considering in specific circumstances.

- Your work involves highly sensitive client data that you cannot risk sending to a third party**  
Mental health records, immigration case details, domestic violence documentation, financial information for vulnerable clients. If the consequence of a vendor data breach would be catastrophic for the people you serve, local processing eliminates that risk entirely.
- Your regulatory environment prohibits third-party data processing**  
Some funding agreements, state regulations, or sector-specific compliance frameworks restrict where data can be processed and stored. If cloud AI tools are incompatible with those requirements, local is the only compliant path.
- You operate in low-connectivity environments**  
Field offices, rural programs, or disaster response contexts where reliable internet access can't be guaranteed. A local model works offline: a cloud tool doesn't.
- You have the technical capacity to set it up and maintain it**  
Local AI is not a burden to take on lightly. If your organization has IT staff or a technically capable staff member who can own the setup, maintenance, and security patching, it's viable. Without that capacity, it creates more risk than it removes.

## GETTING STARTED WITH LOCAL AI **The practical path if you decide to explore it**

If the use case is right and the hardware is available, local AI is more accessible than it was two years ago. A few tools have made the setup significantly less technical.

- ❑ **Ollama is the most accessible starting point**

A free, open-source tool that handles downloading, running, and managing local AI models with minimal technical setup. Works on Mac, Windows, and Linux. Start at [ollama.com](https://ollama.com). A staff member comfortable with installing software can get a model running in under an hour.

- ❑ **LM Studio offers a graphical interface**

For organizations where command-line tools are a barrier, LM Studio provides a desktop application for downloading and running local models without any coding. Free for personal and nonprofit use. [lmstudio.ai](https://lmstudio.ai).

- ❑ **Start with a smaller model to test your hardware**

Models are available in different sizes. Start with a smaller one (7B parameters or less) to confirm your hardware can run it before investing time in a larger, more capable model. Ollama's model library lists hardware requirements for each option.

- ❑ **Local models still require the same oversight practices**

Running AI locally solves the data transit problem. It does not solve the hallucination problem, the bias problem, or the human oversight requirement. Every output still needs a human read before it's used.

**The honest answer for most nonprofits:** Start with a cloud tool that has strong data privacy terms, a clear no-training policy, and SOC 2 certification. That covers the vast majority of use cases. Revisit local models when and if you encounter a specific use case where the data sensitivity is high enough that third-party processing is genuinely unacceptable: not as a general preference, but as a real compliance or safety requirement.

## Minnesota-specific resources

---

Minnesota has a strong nonprofit infrastructure and a handful of resources to have on your radar before you start. These are specific to the state's ecosystem and regulatory environment.

SECTOR SUPPORT **Organizations that help Minnesota nonprofits with technology**

---

- Minnesota Council of Nonprofits (MNCNonprofits.org)**

The state's primary nonprofit membership and advocacy organization. Publishes guides, hosts trainings, and maintains a salary and benefits survey used across the sector. Check their Technology Resources section for current AI guidance and peer connections.
- MAP for Nonprofits (mapfornonprofits.org)**

Provides consulting, capacity building, and technology planning support specifically for Minnesota nonprofits. A good first call if your organization needs hands-on help with a technology adoption plan and doesn't know where to start.
- Propel Nonprofits (propelnonprofits.org)**

Minneapolis based. Provides financial consulting, lending, and strategic capacity building for nonprofits across Minnesota and the surrounding region. Useful when you are deciding whether an AI subscription is worth a multi-year commitment or whether a hardware investment fits your budget. Their finance team is well known in the sector and their lending arm can bridge gaps that a grant cycle cannot.
- TechSoup (techsoup.org)**

The national clearinghouse for verified nonprofit discounts on software and technology, including AI tools. Minnesota nonprofits can verify their eligibility and access discounted pricing for Microsoft, Google, Canva, and others through TechSoup before budgeting any tool at full price.
- Charities Review Council (smartgivers.org)**

Minnesota based accountability and standards organization. Publishes the Meets Standards certification that many Minnesota funders look for when evaluating grantees. Increasingly involved in AI disclosure standards for nonprofits. If you are building a public AI use policy, check their current guidance first.
- Initiative Foundations of Minnesota**

Six regional foundations covering greater Minnesota (Initiative Foundation, Northland, Northwest, Southern, Southwest, West Central). If your organization is outside the Twin Cities metro, your regional initiative foundation is often the best first call for technology planning support, small operating grants, and peer connections. Several have been funding AI readiness conversations for rural and small-town nonprofits since 2024.

## REGULATORY CONTEXT **What Minnesota law adds to your data obligations**

### Minnesota Government Data Practices Act (MGDPA)

If your organization receives state funding or contracts with government agencies, you may be subject to the MGDPA, which governs how government data is classified, handled, and disclosed. AI tools that process government-classified data introduce additional compliance requirements. If this applies to you, confirm your AI data policy addresses MGDPA requirements before any tool goes live. Your legal counsel or the Minnesota Department of Administration can advise.

### Minnesota's privacy laws have evolved

Minnesota enacted a consumer data privacy law in 2023 (effective July 2025) that includes provisions affecting how organizations handle personal data. Nonprofits with specific exemptions should verify whether those exemptions apply to their AI use cases. Get legal guidance before you build workflows around sensitive data.

### Funder-specific data requirements

Several major Minnesota funders, including the Otto Bremer Trust, Bush Foundation, and McKnight Foundation, are developing or have published responsible AI frameworks for grantees. Before submitting to any of these funders, check their current guidance on AI use in proposals and program delivery. Proactively disclosing your AI policy can strengthen a relationship with funders who are actively thinking through these questions.

## PEER LEARNING **You're not the only one figuring this out**

### Minnesota Emerging Leaders in Philanthropy and Nonprofit Management

Peer networks within Minnesota's nonprofit sector are increasingly hosting conversations about AI adoption. If your organization is a member of MNCNonprofits or connected to a regional United Way, ask whether peer roundtables on technology adoption are available. Learning from an organization that already worked through the same decisions is faster than starting from scratch.

### The Minnesota Council of Foundations (mcf.org)

Tracks what major Minnesota funders are thinking and doing on AI. Useful if you want to understand the funder perspective before your next proposal cycle.

**Important:** Minnesota's nonprofit sector is dense and well-connected. If you're a small organization without in-house technical expertise, the resources above exist to help you, not to sell you something. Use them.

# AI and grant writing

Grant writing is where most nonprofits see the clearest return from AI. It's also where the risks are specific enough to address directly. The time savings are real. So are the ways it can go wrong.

## WHAT AI DOES WELL HERE **The tasks worth offloading**

- Drafting narrative sections from your program data**

Feed the AI your outcomes, population served, program model, and budget context. It produces a structured first draft significantly faster than starting from a blank page. Your job is to make it accurate and specific: AI drafts generalize.
- Adapting existing narratives for new funders**

You've written this program description twenty times for twenty funders. AI can reframe your existing narrative for a new funder's priorities, tone, and word limits in minutes. This is one of the most time-intensive parts of grant work and one of the safest AI applications.
- Summarizing program reports for funder updates**

Paste in your data, service numbers, and staff notes. AI produces a readable summary. Staff edit for accuracy and voice. The summary is a starting point, not a finished product.
- Proofreading, tightening, and formatting**

AI catches grammatical errors, awkward phrasing, and inconsistencies faster than a tired grant writer at the end of a deadline week. Use it as a final pass, not a replacement for review.

## THE RISKS THAT MATTER **What can go wrong and how to prevent it**

---

### AI invents plausible-sounding statistics

This is the most dangerous failure mode. AI will generate specific impact numbers, census data, research citations, or program outcomes that sound credible and are simply wrong. Every statistic in an AI-assisted grant narrative must be verified against your actual data before submission. No exceptions.

### Some funders prohibit or restrict AI-assisted proposals

Several major foundations have begun requiring disclosure of AI use in grant applications, and a smaller number prohibit it entirely. Check the funder's current guidelines before using AI to draft a proposal. This is especially relevant for Minnesota funders who are actively developing AI policies: see the MN resources section.

### AI doesn't know your community

AI drafts substitute generalized descriptions of "underserved communities" or "at-risk populations" for the specific, grounded language that comes from actually knowing the people you serve. Funders who know your community notice. Edit relentlessly for specificity.

### Misrepresentation is a compliance risk regardless of how it happened

If an AI-assisted narrative misrepresents your program outcomes, your staff size, your service numbers, or your use of funds, that is still a misrepresentation. The tool that produced it doesn't change your organization's accountability.

## RECOMMENDED WORKFLOW **How to use AI in grant writing without the risks**

- Start with your real data, not a blank prompt**  
Give the AI your actual numbers, actual program description, actual population, and actual outcomes before asking it to draft. The quality of the output depends almost entirely on the quality of what you put in.
- Treat every output as a first draft that needs a fact-checker**  
Have someone other than the AI user read the draft specifically for invented details. Build this into your review process explicitly: not as an afterthought.
- Keep the funder's guidelines open while editing**  
AI drafts don't read RFPs. You do. Make sure the final narrative answers the questions actually asked, not a generalized version of what the AI thought was being asked.
- Document your process for funders who ask**  
Some funders will ask whether AI was used. Know your answer before they ask it. "We used AI to produce an initial draft, which was reviewed, fact-checked, and substantially edited by staff" is honest and defensible.

**The honest summary:** AI can cut grant writing time significantly, most of that time savings comes from drafting and reformatting tasks that don't require human judgment. The tasks that do require judgment: knowing your community, verifying your data, understanding what a funder actually cares about: still require a person. Use AI to clear the low-value work so your grant writers have more time for the high-value work.



### ABOUT THE AUTHOR

**Mandy Hathaway**

**AI Ethics Consultant & Technology Specialist · MA in Ethical Technology & AI**

Mandy Hathaway is an AI ethics consultant and technology specialist with 30+ years of technology experience and an MA in Ethical Technology & Artificial Intelligence from Metropolitan State University. She holds a certificate in AI for Nonprofits from Microsoft Elevate, LinkedIn, and NetHope, and has published on AI bias and responsible AI development. She works with nonprofits and mission-driven organizations on AI adoption: ethics, policy, tool selection, and staff training.

NEED MORE THAN THIS GUIDE?

## Let's talk through it together

Every organization is different. If you've worked through this guide and still have questions about a specific tool, a tricky data situation, or where to start, reach out. I work with nonprofits on exactly these questions.

**WEBSITE**

[mandyhathaway.com](http://mandyhathaway.com)

**EMAIL**

[hathaway.mandyl@gmail.com](mailto:hathaway.mandyl@gmail.com)

**LINKEDIN**

[linkedin.com/in/mandyhathaway](https://www.linkedin.com/in/mandyhathaway)

**SCHEDULE A CALL**

[calendly.com/mandyhathaway](https://calendly.com/mandyhathaway)

© 2026 Mandy Hathaway · [mandyhathaway.com](http://mandyhathaway.com) · AI Ethics Consultant & Technology Specialist