WHITE PAPER

AI governance for a small organization

Your staff are already using AI. The question is whether the organization has decided anything about it. A proportionate approach for associations and nonprofits of 10 to 80 people, ending in a one-page AI policy your board can read in five minutes and approve.

01 Why governance, and why small

02 What AI may see

03 Who reviews what

04 Telling members and donors

05 Measuring it, and switching it off

06 Ten questions for every vendor

07 The one-page AI policy

08 From approval to practice

For
Anyone deciding how their organization uses AI

Time to read
About 25 minutes

Use it
Before the next AI tool you buy

A NOTE BEFORE YOU START

Governance in a small organization is not a committee. It is a handful of decisions, written down.

In most organizations I visit, AI arrived without anyone deciding it should. A development officer drafts appeal letters with a chatbot. Someone in membership pastes a renewal list into a tool to tidy the names. A note-taker joins board calls because it came with the video subscription.

None of that is reckless. It is people trying to do more with less, which is what small teams have always done. But it means the organization is already making choices about member and donor information, and nobody made them on purpose.

The governance frameworks written for large companies do not help much here. They assume a risk team, a legal department and an inventory of models. A 30 person association has an executive director, a board that meets quarterly and a part-time IT contractor. What it needs is proportionate: a few clear rules, one owner, a short list of what is allowed, and a way to stop.

This paper is the version I would want if I were running a small organization. It ends in a one-page policy you can adapt and take to your board. It is not legal advice, and where a lawyer should look, I say so plainly.

Sunil Dua
Executive Director and Principal Consultant, Kloudserv

The short version

- Find out what is already in use before writing any rules.
- Sort information into four tiers and say which tiers AI may touch.
- Match review to risk. Decisions about people stay with people.
- Tell members and donors plainly when AI is part of what they receive.
- Measure a few things, and name who can switch it off.
- Ask every vendor the same ten questions.

What this is not

A legal opinion, or a reason to slow down. Good rules make it easier to say yes, because people know where the lines are.

Worth remembering. The point of governance is not to stop staff using AI. It is to make sure that when a member or donor asks "how did you use my information?", the organization has an honest answer it chose in advance.

01 · WHY GOVERNANCE, AND WHY SMALL

The risk is rarely the technology. It is the decisions nobody made.

Large organizations govern AI because of scale: thousands of staff, many systems, regulators watching. Small organizations face a different shape of risk. Fewer people, so fewer second pairs of eyes. Systems run by one or two staff. Information that is personal and often sensitive: certification records, giving histories, scholarship applications, volunteer background checks. And a reputation that rests on trust far more than on brand.

When something goes wrong in a small organization, it is rarely dramatic. It is a donor list pasted into a personal account whose terms allow inputs to be kept. A member question about CE credits answered with a confident, wrong statement. A grant report with a figure nobody can trace back. Each one is small. Together they wear down the thing a membership or donor organization cannot buy back.

Three tests of proportion

Could this affect a person? If an output could change someone's membership, credit, grant, scholarship, job or standing, it needs a human decision and a record.

Does it touch personal information? If it reads member, donor, staff or volunteer data, the tool's terms matter more than its features.

Would we explain it happily? If you would hesitate to describe the use to a member, a donor or a funder, that hesitation is the signal.

Start with what is already happening

Before writing any rule, find out what is in use. Not as an audit, and not to catch anyone out. Ask each team three questions, in a staff meeting or a short form: which AI tools do you use, for what, and what information goes into them. Include the features that arrived inside software you already pay for: meeting note-takers, writing assistants in email and documents, AI features in your CRM or your fundraising and event platforms.

In most small organizations the list is longer than leadership expects, and almost all of it is reasonable. That list becomes your first register of uses. It tells you where the rules need to land, and it means the policy you write governs what is really happening rather than what you imagine.

Make it safe to answer. If people think the survey is the first step to a ban, they will under-report, and you will govern a picture that is not real. Say up front that the aim is to approve what is working and agree a few lines, not to take tools away.

02 · WHAT AI MAY SEE

Sort your information into four tiers, then decide what each tier may touch.

Most AI risk in a small organization comes down to one question: what information went into which tool, under whose terms. The workable answer is a short classification every staff member can hold in their head.

TIER	WHAT IT INCLUDES	EXAMPLES	AI USE
1 · Public	Anything published or meant to be	Website copy, published reports, event descriptions, job postings	Any approved tool
2 · Internal	Work material not for the public and not about individuals	Draft plans, budgets, policies, agendas, process notes	Approved tools under business terms that do not train on your data
3 · Personal	Information about identifiable people	Member and donor records, contact lists, giving and attendance history, staff files	Only in approved systems under contract, for a purpose recorded in the register
4 · Sensitive	Personal information where exposure could cause real harm	Health or disability details, financial hardship, background checks, anything about minors, scholarship and grant applications, complaints and discipline	Not in AI tools unless the executive director approves the specific use in writing, after legal review

Two distinctions matter more than any others. The first is between the consumer and business versions of the same tool. A personal account often comes with terms that allow your inputs to be kept or used to improve the service; a business agreement usually does not. Same screen, very different promise. Check the terms you actually accepted, not the product page.

The second is between a separate tool and a system you already control. An AI feature inside your CRM, working on records already held there, under a contract you have read, is a different risk from copying the same records into a website. Keeping data where it already lives is usually the safer default.

Practical rules that follow

- Keep an approved tools list on one page, with the account type for each.
- Personal accounts are for Tier 1 only.
- Remove names and identifiers when the task does not need them. Most drafting does not.
- Nobody uploads a full export of members or donors to any AI tool without a named approval.

Legal review is advisable here. If you hold health information, information about children, or members and donors in several states, provinces or countries, privacy law, funder terms and your existing contracts may set rules that override anything in this paper. Ask counsel to read your four tiers once, before the board sees them.

03 · WHO REVIEWS WHAT

Match the review to the risk. Decisions about people stay with people.

"A human reviews everything" sounds safe and fails in practice. Reviewing every AI-assisted social post slows the team down and teaches people to skim. Reviewing nothing about a scholarship shortlist is how an organization ends up defending a decision it cannot explain. Three levels work better than one rule.

LEVEL	TYPICAL USES	WHO REVIEWS	WHAT IS RECORDED
Light	Internal drafting, summarizing your own documents, brainstorming, formatting, first drafts for colleagues	The person using it, as they would check any of their own work	Nothing beyond normal practice
Standard	Anything sent to members, donors, funders or the public. Answers to member questions. Figures in reports and grant applications	A second person, or the author against a named source, before it goes out	That AI was used, and who approved the final version
Decision	Eligibility, membership status and lapses, CE credits and certification, grants, scholarships, awards, hiring, complaints, anything that changes someone's standing	A named staff member decides. AI may sort, summarize or flag. It never decides	The decision, the reason in plain English, and the person who made it

Name one owner

Every policy needs one person who owns it. In a small organization that is usually the COO or director of operations, sometimes the executive director. Not IT alone: the hard questions are about members, donors and fairness, not configuration. IT advises; operations owns.

The owner keeps the register and the approved tools list, approves new Standard level uses, and brings any new Decision level use to the executive director, or to the board if the board has kept that decision for itself.

Make the register real

The register of uses is a plain list: the use, the tool, the tier of information, the review level, the owner, the date approved, and how to switch it off. A spreadsheet is fine.

Its value shows up the first time a board member asks "where are we using AI?" and the answer takes a minute rather than a week. It is also what makes disclosure honest, because you cannot describe uses you have not written down.

A test worth running. If the AI tool were unavailable tomorrow, could the person who owns each Decision level process still make every decision, and explain it? If not, the tool is already deciding. It just has not been given the title.

04 · TELLING MEMBERS AND DONORS

Tell people when AI is part of what they receive. Plainly, and before they have to ask.

Members and donors do not expect you to avoid AI. They do expect not to be surprised. When an organization loses trust over AI, it is usually over discovery rather than use: a member learns that the warm, personal reply was generated, or a donor finds out their giving history went through a tool they had never heard of. Disclosure has three parts.

In your privacy notice. A short paragraph: you use AI tools to help staff work, what kinds of information they may process, that you choose tools under business terms, and that people make decisions about members and donors.

At the point of contact. Where someone deals with AI directly, such as a website assistant or an automated first reply, say so in the interface. One sentence is enough.

On request. Anyone who asks how AI was used with their information gets a straight answer from a named person, drawn from the register.

Wording you can adapt

PRIVACY NOTICE

"We use AI tools to help our staff draft, summarize and organize their work. Where these tools process information about members or donors, we use services under business agreements that do not allow them to train on your information. Decisions about membership, credits, grants and gifts are made by people. To find out more, contact [name] at [email]."

WEBSITE ASSISTANT

"This assistant uses AI to answer common questions and can make mistakes. For anything about your membership, credits or account, a member of our staff will confirm."

Where disclosure is not needed

You do not need to label every email a staff member drafted with help, any more than you label spell check. The test is whether a reasonable member or donor would feel misled if they found out. A staff member tidying their own letter: no. An appeal signed by the executive director but written and personalized by a system and sent to thousands: say something, or at least make sure your privacy notice covers it honestly.

Funders may have their own expectations about AI in applications and reports. Read the grant terms, and when they are silent and the use is material, ask the program officer. A short question now costs less than an awkward one later.

Have counsel read the privacy notice change. Privacy notices can carry legal weight, and the right wording depends on where your members and donors live. The paragraph above is a starting point, not a finished notice.

05 · MEASURING IT, AND SWITCHING IT OFF

If you cannot say whether it is working, and cannot stop it, you are not governing it.

Every Standard or Decision level use should have a reason that can be checked. Not a dashboard: four or five things, looked at each quarter by the owner, and summarized for the board once a year.

WHAT TO MEASURE	WHY IT MATTERS	HOW, SIMPLY
Time saved, or work done that would not otherwise happen	The reason you adopted it	Ask the people using it each quarter for a rough estimate. Honest beats precise.
Errors caught in review	Shows whether the review level is right	A tally kept by the reviewer: none, minor, serious
Errors that reached a member or donor	The number that matters most	Logged like any complaint, with what happened next
Questions or complaints about AI	An early warning on trust	Tagged in your shared inbox or CRM
Cost, including review time	Whether it still pays its way	Subscription plus staff time spent checking, each quarter

The switch-off

Every approved use needs a way to stop it within a day, and a person with the authority to do so without asking anyone. Write down three things for each Standard and Decision level use.

Who can switch it off. The policy owner and the executive director, each acting alone. Nobody needs a meeting to stop a problem.

How. The setting, account or integration to disable, written down in plain English, not held in one person's memory.

What happens instead. The manual process the team falls back to, so stopping the tool does not stop the work.

Good reasons to switch off include a serious error reaching a member or donor, a vendor changing its terms or being acquired, a security incident at the vendor, or a measure moving the wrong way two quarters running. Switching off is not a failure. It is the control working as intended.

Test it once. Turning a tool off for a day, on purpose, tells you more about how much you depend on it than any risk register will.

06 · TEN QUESTIONS FOR EVERY VENDOR

Ask before any AI tool touches personal information.

Ask in writing, keep the answers with the contract, and ask again at renewal. A vendor that cannot answer clearly has still given you an answer.

	QUESTION	A GOOD ANSWER SOUNDS LIKE
1	Will our data be used to train or improve your models, or anyone else's?	No, under our business terms, and here is the clause that says so.
2	Where is our data stored and processed, and by which other companies?	Named locations and a published list of subprocessors.
3	How long do you keep what we put in and what comes out, and can we delete it?	A stated period, and deletion on request with confirmation.
4	Who at your company can see our data, and when?	A small set of roles, for support only, with access logged.
5	What happens to our data if we leave, or if you are acquired?	Export in a usable format, deletion afterwards, and terms that survive a sale.
6	Which AI models sit behind the product, and will you tell us when that changes?	Named providers, and notice before a material change.
7	How do we turn the AI features off, for one person or for everyone?	An admin setting they can show you on the call.
8	What independent security assurance can you share?	A recent independent audit or assessment report, shared under a confidentiality agreement.
9	Will you sign our data protection terms, or show us yours?	Yes, and here is our standard agreement to review.
10	What do you record about how AI was used in our account, and can we see it?	Usage logs your administrator can view or export.

Ask the same questions about AI features inside software you already use. Features are often switched on by default in an update, and the terms that cover them may be new, sitting in a document nobody on your team has read. Your CRM, email platform, event system and video tool are all worth a look.

None of this needs to be adversarial. Good vendors expect these questions and answer them quickly. The slow, vague answer is the one to notice.

Where legal review earns its fee. If a vendor will process Tier 3 or Tier 4 information, have counsel read the agreement and the data protection terms before you sign. That is a short, specific piece of work, and it is the part of AI governance where a lawyer's eye matters most.

07 · COPY THIS

The one-page AI policy

Adapt the bracketed parts and keep it to one page. Have counsel review it before it goes to the board.

[Organization] AI Use Policy

Approved by the Board on [date] · Owner: [role]
Reviewed every twelve months, or sooner if uses change materially

1. Purpose. We use AI to help our staff serve members, donors and our mission with less wasted effort. This policy sets out how we do that while protecting the people whose information we hold and the trust they place in us.

2. Scope. It applies to all staff, contractors and volunteers who use AI tools for our work, including AI features inside software we already use.

3. Approved tools. The policy owner keeps a list of approved tools and the account type for each. Only approved tools may be used with information beyond Tier 1. Personal accounts may be used with public information only.

4. What AI may see. Our information falls into four tiers. Tier 1, Public: any approved tool. Tier 2, Internal: approved tools under business terms that do not train on our data. Tier 3, Personal: approved systems under contract, for a purpose recorded in the register. Tier 4, Sensitive: only with the written approval of the Executive Director, after legal review.

5. Human review. Every use has a review level. Light: the user checks their own work. Standard: anything sent to members, donors, funders or the public is checked by a person before it goes out. Decision: AI may assist but never decides matters affecting a person's membership, credits, certification, grants, scholarships, employment or standing. A named person makes and records each such decision and its reason.

6. Disclosure. Our privacy notice explains how we use AI. We say so wherever people interact with AI directly. Anyone may ask how AI was used with their information and will receive a clear answer.

7. Register. The policy owner keeps a register of approved uses: purpose, tool, tier, review level, owner, date approved and how to switch it off.

8. Measuring. For each Standard and Decision level use, the owner reviews time saved, errors, complaints and cost every quarter.

9. Switching off. The policy owner or the Executive Director may suspend any AI use at once, without further approval, after a serious error, a security concern, a change in vendor terms or a loss of confidence. Every use has a written fallback.

10. Vendors. Before any tool processes Tier 3 or Tier 4 information, the owner obtains written answers to our vendor questions, and counsel reviews the terms where advised.

11. Incidents. Anyone who believes personal information went into an unapproved tool, or that an AI error reached a member or donor, tells the policy owner the same day. Reporting in good faith is never treated as misconduct.

12. Board oversight. The board approves this policy and [any new Decision level use]. The owner reports to the board [annually] on the register, incidents, measures and proposed changes.

Board Chair	Executive Director	Date approved

Attach the approved tools list and the register to the board paper. They change more often than the policy, so they sit with management, and the policy stays stable.

08 · FROM APPROVAL TO PRACTICE

The board approves the policy. The staff decide whether it is real.

Boards in small organizations are right to be wary of policies that live in a binder. What makes an AI policy real is not its wording. It is whether the approved tools list is current, whether staff know the four tiers without looking, and whether the first incident is handled calmly enough that the second one gets reported too.

A 90 day plan

Weeks 1 to 3	Ask every team what they use. Build the first register from the answers. Agree who owns the policy.
Weeks 4 to 6	Draft the four tiers and the approved tools list. Send the ten questions to the vendors behind any Tier 3 use. Ask counsel to read the tiers and the main vendor terms.
Weeks 7 to 9	Take the one-page policy to the board or its executive committee, with the register attached. Update the privacy notice once approved.
Weeks 10 to 13	Run a thirty minute staff session using examples from your own register. Test one switch-off on purpose. Collect the first quarter's measures.

What to bring to the board

The one-page policy, the register as it stands (even if it is short), a note on what legal review covered, and one question for the board to decide: which uses, if any, it wants to approve itself rather than delegate. Many boards choose to approve Decision level uses and leave everything else to the executive director. That is a sensible split, and it keeps the board's attention where the risk is.

Four mistakes worth avoiding

A ban. It moves use onto personal accounts, out of sight, which is the exact outcome governance exists to prevent.

A policy borrowed from a large company. Thirty pages nobody reads, naming roles the organization does not have.

IT as the owner. The judgments are about members, donors and fairness. IT advises on the tools; operations owns the decisions.

Approved once, reviewed never. Tools change month to month. A policy without a yearly review describes last year.

WHAT TO DO NEXT

Two steps, this month.

01 Ask your staff what they already use

Three questions, asked in a way that makes an honest answer safe: which tools, for what, and what information goes in. Write the answers down. That is your first register, and most of the policy follows from it.

02 Put the one-page policy on the next board agenda

Adapt page 9, attach the register, and ask counsel to read the tiers and the privacy notice change before the meeting. Ask the board one clear question: approve the policy, and decide which uses it wants to keep for itself.

Everything else in this paper can follow over the next quarter. A small organization does not need to govern AI the way a bank does. It needs to know what it is doing, decide a few lines on purpose, and be able to explain both to the people it serves.

That is within reach of any team, whatever systems it runs.

WHERE THIS CONNECTS TO YOUR SYSTEMS

Governance is easier when you know where your member and donor information lives and who looks after it. If that picture is unclear, the Digital Readiness Blueprint maps what you run, what each system holds and what would need to be in place before AI can be trusted with it.

This paper is general guidance for associations and nonprofits. It is not legal advice. Requirements differ by jurisdiction, sector and contract, so have counsel review your policy, privacy notice and vendor terms before you rely on them.