



Responsible AI Governance Toolkit for Canadian Recreation and Community Service Organizations — Policy Template Reference Guide for Small Organizations / Teams

By: [Toby Nwabuogor, AIGP](#) and [Dr. Julie Booke](#)

Version 1.0 | February 2026



About this resource

This policy template is designed for **small teams or organizations with fewer than 10 members**. It includes the essential governance elements needed for safe, trustworthy AI use, without the overhead of enterprise frameworks that small teams cannot sustain. The core principles can be scaled up or down based on team size and capacity, while recognizing that larger organizations typically have more stakeholders, stricter compliance needs, and more formal approvals.

Further assistance is available here: <https://www.civicplay.ai/ai-governance-toolkit>

Who this is for

- Small nonprofits and recreation organizations
- Teams with limited administrative capacity
- Organizations where one person holds multiple roles
- Groups that need to act quickly while maintaining responsible practices

What makes this different

- **Streamlined structure:** One lead instead of multiple committees, with optional advisors
- **Living document:** Built to evolve through community learning and scheduled reviews
- **Practical focus:** Implementable in 48 hours to 1 week
- **Learning culture:** Prioritizes questions and improvement over compliance and punishment

Also works for

- Small teams or departments within larger organizations (for example, a 5-person recreation team within a municipal government)
- Departments that have limited AI governance capacity even if the broader organization is large

If you are a team within a larger organization

- Check whether your organization already has an AI policy or data classification approach, and align with that first.
- Confirm whether your team can approve tools independently or whether organizational sign-off is required.



- Ensure your traffic light system matches your organization’s data policies.
- Coordinate with IT, legal, or privacy teams before implementing.

Use this reference guide as a practical framework for adopting responsible AI governance, not as a strict code.

The [blank policy template](#) is available separately online for free. The expanded guidance below is intended for the prospective AI Governance Lead, relevant leadership, and anyone who wants a fuller picture of how to apply the policy in real situations.



Reference Guide Table of Contents

Responsible AI Governance Toolkit for Canadian Recreation and Community Service Organizations — Policy Template Reference Guide for Small Organizations / Teams	1
About this resource	2
Who this is for	2
What makes this different	2
Also works for	2
If you are a team within a larger organization	2
Reference Guide Table of Contents	4
How to use this reference guide	7
Who should read what	7
What this guide is for	7
Role guides (what to do, by role)	7
Executive or leadership	7
AI Governance Lead (policy owner)	8
Managers or supervisors	8
Frontline staff, contractors, or volunteers	8
Governance structure (visual)	9
Template guidance notes (read alongside the policy)	10
[Organization Name] Responsible AI Use Policy - The policy template with expanded details	12
Purpose	12
What counts as AI for this policy	12
Mission alignment	13
Guiding principles: the 3Cs framework	14
Accountability	15
Risk-based classification: the traffic light system	16
Human oversight requirements	17
Ongoing learning and sharing	18
Policy review	19
Approved AI tools	19
Contact and questions	19
When to reach out	19
Acknowledgment	20
Implementation guide	21
First 48 hours	21

First 90 days	21
Success indicators	21
Decision rules (fast, consistent answers)	22
The 60-second decision tree	22
Non-negotiables	22
Verification (how to reduce harm)	22
General verification checklist	22
Verification triggers (always double-check)	23
Incident and near-miss handling (no-blame learning)	23
What counts as an incident	23
What to do immediately	23
Common challenges and solutions	23
Challenge: “We’re too small to need a policy”	23
Challenge: “I don’t have time to be the Governance Lead”	23
Challenge: “Our staff won’t follow this”	24
Challenge: “What if someone breaks the policy?”	24
Challenge: “AI is changing too fast for a policy”	24
Challenge: “We can’t afford enterprise AI tools”	24
Sector scenarios (small org, recreation and community services)	25
Scenario A: Youth program newsletter draft	25
Scenario B: Subsidy or financial assistance communication	25
Scenario C: Incident report or safety event summary	25
Legal, privacy, and accessibility (plain-language navigator)	25
When to get help	25
Accessibility basics for AI-assisted communications	26
Maintenance rhythm (keep it sustainable)	26
Monthly (2 to 4 hours)	26
Scheduled reviews	26
Policy implementation checklist	27
Policy foundation	27
Tools and access	27
Support systems	27
Learning culture	27
Ongoing review	27
Disclaimer and disclosure	28
Transparency statement	28
Policy disclaimer	28



6

License & Attribution	29
Authorship and development disclosure	29
Appendix: related resources (optional)	31



How to use this reference guide

Who should read what

- **Executive or leadership:** Read *Purpose*, *Non-negotiables*, *Decision rules*, *Incident response*, and the *Executive role guide*.
- **AI Governance Lead:** Read everything. This is your operating manual.
- **Managers or supervisors:** Read the *Manager role guide*, the decision tree, and the verification checklists.
- **Staff, contractors, or volunteers:** Read the frontline guidance and keep the quick-reference rules handy.

What this guide is for

- Answering: “Is AI appropriate here?”
- Creating consistency: “If two people ask the same question, do we give the same answer?”
- Keeping the policy alive: “How do we update guidance without adding bureaucracy?”

Role guides (what to do, by role)

AI governance is not an extra task or a one-person job. It is how we protect our mission and the trust our community places in us while we use new tools to work faster and better. Everyone has a role: leadership sets expectations and supports responsible choices, the AI Governance Lead provides clarity and keeps guidance current, managers coach and approve higher-impact uses, and staff bring good judgment to daily work by checking the data, verifying outputs, and speaking up early when something does not feel right. When we treat responsible AI use as part of the mission, we build a culture where learning is shared, risks are caught early, and people stay accountable for the outcomes.

Executive or leadership

Your job is to: set direction, protect trust, and ensure the policy is used.

- Name the AI Governance Lead and support their decision-making.
- Approve the policy and any major updates.
- Model the behaviour: ask “What data is this?” and “Can we verify this?”



- Treat incidents as **learning events first**, unless intent or harm requires escalation.

AI Governance Lead (policy owner)

Your job is to: be the point of clarity, not a bottleneck.

- Maintain the approved tools list and the “what’s allowed” guidance.
- Answer questions within a reasonable timeframe. Set a clear response standard.
- Capture recurring questions and convert them into guidance.
- Maintain a lightweight incident and near-miss log.
- Run scheduled reviews and publish changes.

Managers or supervisors

Your job is to: approve, review, and coach.

Use this when someone asks, “Can I use AI for this?”

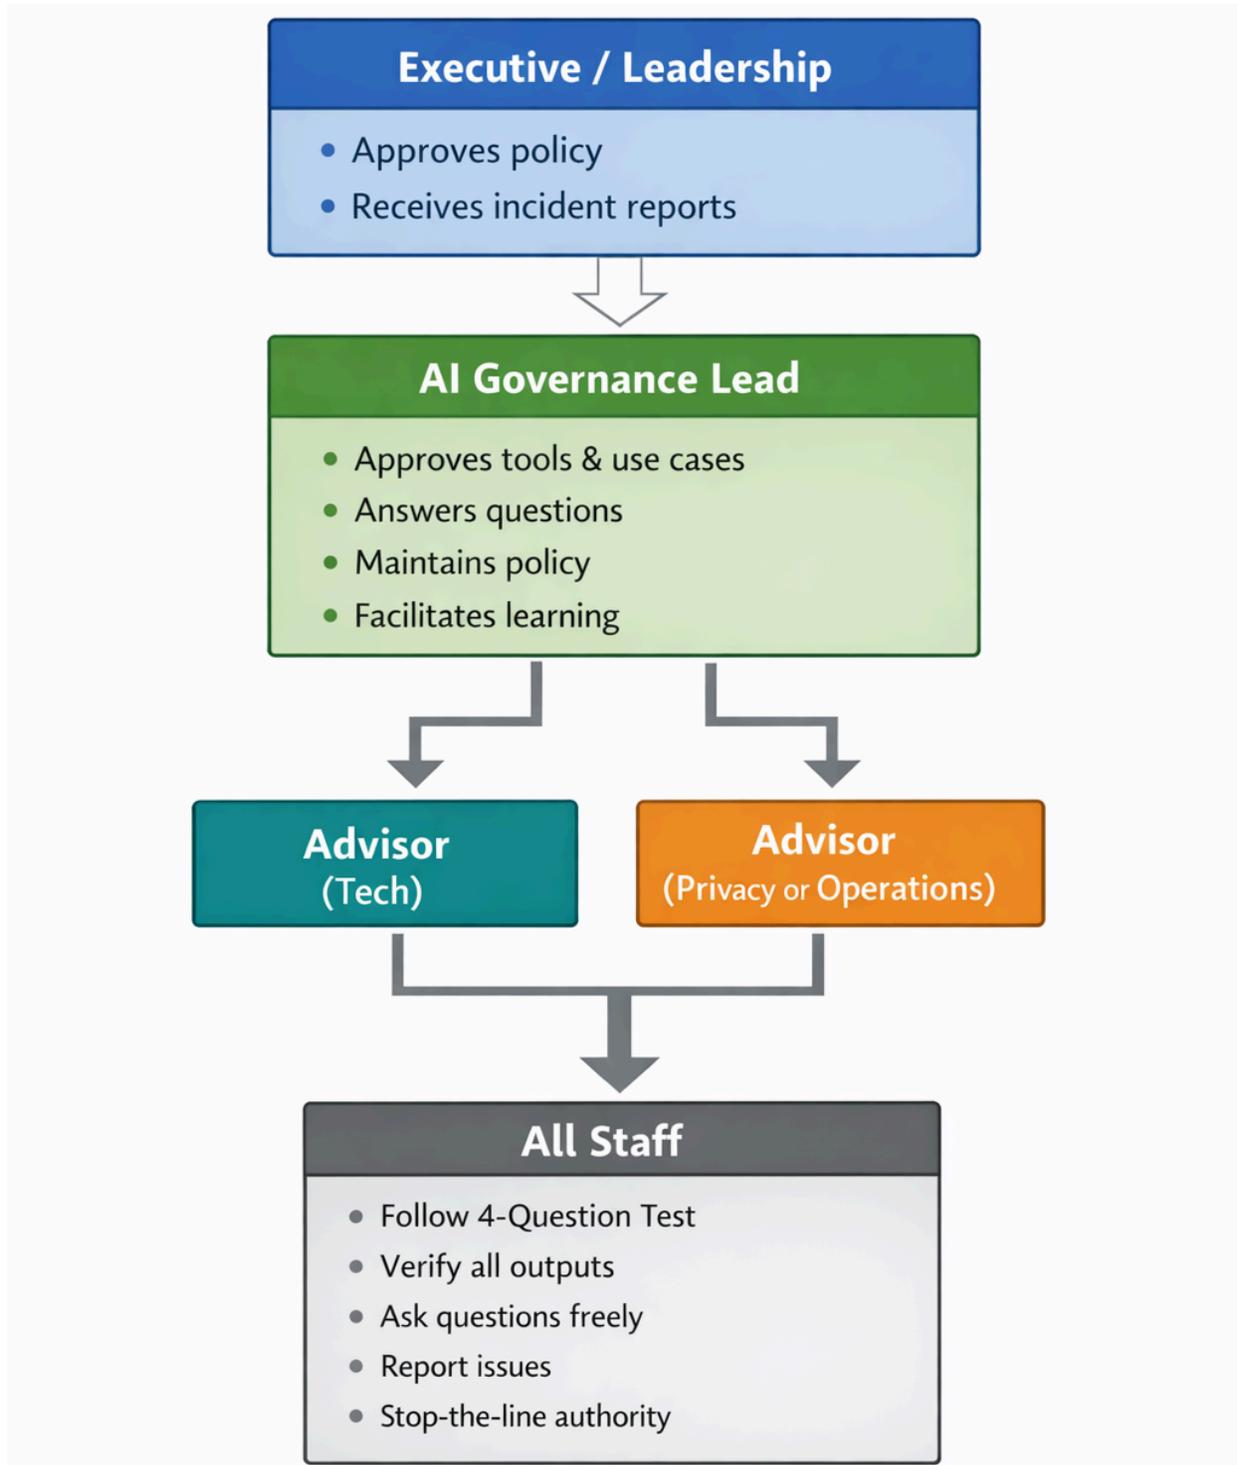
- What is the outcome, and who will see it?
- What data color is involved?
- Can the output be verified?
- What is the harm if it is wrong?
- Does this cross a non-negotiable boundary?

Frontline staff, contractors, or volunteers

Your job is to: follow the checks, ask early, and report problems.

- Use the 4-Question Test.
- Do not paste sensitive information.
- Review every output.
- When unsure, ask before sharing externally.

Governance structure



Template guidance notes (read alongside the policy)

- **Purpose** — how to use this policy day to day
 - This policy is for answering: “Is AI appropriate here?” and “What safeguards do we need?”
 - Aim for consistency. If two staff ask the same question, the answer should be the same.
- **What counts as AI** — avoid the “we didn’t know that feature was AI” trap
 - Assume AI can be embedded in tools you already use.
 - When a vendor adds AI features, treat it as “new AI use” until the AI Governance Lead confirms it is safe for the data types involved.
- **Mission alignment** — pick your non-negotiables
 - Write 2 to 4 non-negotiables in plain language.
 - Use them as escalation triggers when a use case feels ethically “off,” even if it seems efficient.
- **3Cs** — practical translation
 - **Care:** Identify who could be harmed if the output is wrong or leaked.
 - **Clarity:** Be able to explain the choice to a community member in one minute.
 - **Control:** Name the human decision-maker for each output.
- **Accountability** — keep the Lead a point of clarity, not a bottleneck
 - Set an expectation for questions (for example, “within 1 business day”).
 - Capture repeat questions and turn them into a short FAQ or guidance note.
- **Traffic Light System** — quick classification and safer anonymization
 - When uncertain, treat data as higher risk.
 - **Never paste personal or confidential information into AI tools.**
 - If you need help drafting something sensitive, **write it without identifiers** (Participant A, age range, no address), then **add the real details back manually** in your own document.
 - If someone could still figure out who it is from context, **treat it as Restricted (Red)** and do not use AI.
 - “Anonymized” still requires caution. If re-identification is plausible, treat it as restricted.
- **Decision rules** — the 60-second check
 - What data is involved? (Green, Yellow, Red)
 - What is the task? Drafting, summarizing, deciding, or communicating?
 - Can you verify the output? If no, stop.
 - Who approves? Make the approver explicit.
 - Would stakeholders trust this if they knew? If no, escalate.
- **Verification** — what to double-check every time



- Facts, dates, costs, names, and numbers.
 - Anything that sounds overly confident.
 - Any legal, policy, or safety claim.
 - Any content involving youth, eligibility, finances, or incidents.
 - **Incident and near-miss handling** — a no-blame learning loop
 - Stop using the output.
 - Notify the AI Governance Lead (and leadership if serious).
 - Record what happened and how it was caught.
 - Update guidance so the same mistake is less likely next time.
 - **Maintenance rhythm** — keep it sustainable
 - Monthly: review questions, near-misses, and tool list updates.
 - Scheduled reviews: follow the cadence in the policy, plus review after major changes or incidents.
-



[Organization Name] Responsible AI Use Policy - The policy template with expanded details

Effective date: *[Date]*

AI Governance Lead (policy owner): *[Name and contact]*

Purpose

This policy enables our team to use AI responsibly while protecting our community's trust. It provides clear guidance on safe AI use, establishes accountability, and creates a learning environment where we can innovate within appropriate guardrails.

What counts as AI for this policy

This policy applies to any system or tool that uses artificial intelligence, machine learning, or automated decision-making to generate content, make predictions, provide recommendations, or take actions.

This includes:

- Text generation tools (ChatGPT, Gemini, Copilot, Claude, etc)
- AI assistants and agents (including autonomous or semi-autonomous agents)
- AI features embedded in software you already use (Microsoft 365 Copilot, Notion AI, Grammarly, etc)
- Image, audio, or video generation tools
- Automated analysis, classification, or recommendation systems
- Any tool that learns from data to produce outputs

Key principle: If a tool uses AI to help you work, this policy applies, even if AI is only one feature in a larger system. When in doubt, ask the AI Governance Lead.



Mission alignment

Before adopting this policy, connect it to what matters most to your organization.

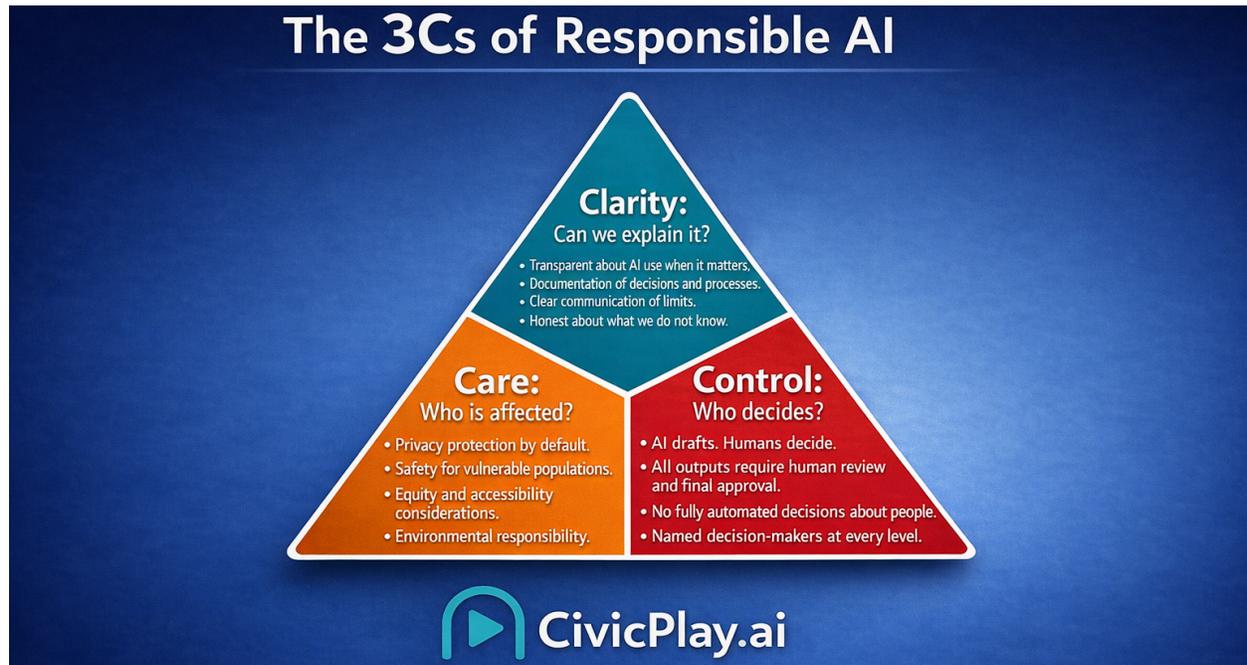
Three key questions:

1. **Our mission:** *[Your organization's mission statement]*
2. **How AI helps:** *AI should help us [specific goal] by [specific contribution]*
3. **What we will not compromise:** *[Non-negotiables]*

Sector-specific considerations (choose what applies):

- **Youth-serving:** No AI tools with age restrictions. Staff supervision required.
 - **Recreation or Sports:** No AI decisions about eligibility or participant readiness.
 - **Social services:** No AI determination of who qualifies for support.
 - **Indigenous-led:** Honor Indigenous data sovereignty and OCAP® principles.
 - **Arts or Cultural:** Protect artist attribution. AI does not replace human creativity in core programming.
-

Guiding principles: the 3Cs framework



Care: Who is affected?

- Privacy protection by default.
- Safety for vulnerable populations.
- Equity and accessibility considerations.
- Environmental responsibility.

Clarity: Can we explain it?

- Transparent about AI use when it matters.
- Documentation of decisions and processes.
- Clear communication of limits.
- Honest about what we do not know.

Control: Who decides?

- AI drafts. Humans decide.
- All outputs require human review and final approval.
- No fully automated decisions about people.
- Named decision-makers at every level.



Accountability

AI Governance Lead: *[Name]*

Responsibilities:

- Primary contact for AI-related questions and concerns.
- Approves new tools and use cases, or coordinates approvals for medium or high-risk uses.
- Maintains this policy and supporting guidance.
- Facilitates learning and policy reviews.

Note: In small organizations, the AI Governance Lead may also be the policy owner with executive accountability. These can be the same person, or split based on your structure.

Need a second opinion? When facing uncertain decisions, the AI Governance Lead should loop in 1 to 2 colleagues with different perspectives rather than deciding alone.

Stop-the-line authority: Anyone can and should stop AI use if they identify a risk. Speak up immediately to your manager or the AI Governance Lead. If you discover a pattern of bias or harmful outputs from an AI tool, flag it immediately to the AI Governance Lead for tool review.



Risk-based classification: the traffic light system

● Green data (Open)

- **What:** Publicly available information, general knowledge, non-confidential content
- **Examples:** Published reports, public newsletters, general program descriptions
- **Tools allowed:** Any approved AI tool
- **Rule:** There is no privacy risk if the information is already public

● Yellow data (Internal)

- **What:** Operational data, internal policies, draft documents, organizational information
- **Examples:** Budget drafts, internal meeting notes, strategic plans
- **Tools allowed:** Enterprise AI tools only
- **Rule:** Never use free or personal AI tools with Yellow data

● Red data (Restricted)

- **What:** Personally identifying information, confidential data, participant details, sensitive organizational information
- **Examples:** Participant contact information, children's information, financial assistance details, medical information, subsidy applications, employee personal files
- **Tools allowed:** Prohibited in cloud AI tools unless an explicit exception is approved through a privacy or legal review
- **Rule:** When in doubt, treat data as Red

Explicit guidance on uses

You can use AI for (with review):

- Drafting internal documents and emails (review before sending)
- Summarizing your own meeting notes
- Brainstorming ideas and planning
- Formatting and editing your writing
- Translation (verified by someone fluent in the language)
- Drafting program descriptions (manager approval required)

You must get approval for:

- External communications to the public
- Social media posts
- Grant applications
- Anything involving participant data (even anonymized)
- Policy or procedure documents
- New AI tools that are not yet approved

You never use AI for:

- Making final decisions about people (hiring, eligibility, discipline)
- Pasting personal, confidential, or participant information into AI tools
- Anything you cannot verify is accurate
- Tasks where a mistake would cause serious harm
- Replacing human judgment in sensitive situations

Human oversight requirements

Before you start: Pause and ask why you are reaching for AI. What problem are you solving? Is AI the right tool for this task? If you cannot articulate a clear purpose, step back and reconsider.

Critical reminder: AI can hallucinate. AI tools sometimes generate false information with complete confidence.



The 4-Question Test (before using AI for any task):

1. **Can I verify it?** If you cannot check accuracy and appropriateness, do not use AI.
2. **Is it confidential?** Check the traffic light system. Never paste Red data.
3. **Do I approve it?** You must review and approve. AI does not decide for you.
4. **Would stakeholders trust this if they knew?** If not, stop and escalate.

Verification checklist (for important outputs):

- Factually accurate and complete
- Appropriate tone for the audience
- Reflects our organizational values
- Free of harmful bias or stereotypes
- Protects privacy and confidentiality
- Accessible and understandable

Ownership: Any content created using AI during work belongs to [Organization Name]. You remain responsible for its accuracy and appropriateness.

Ongoing learning and sharing

AI technology evolves quickly, so we will learn together. This policy focuses on learning, not blame.

Share what works and what does not:

- Share successes (useful prompts, time-saving techniques, helpful applications)
- Share challenges (near-misses, lessons learned, things that did not work)
- Ask questions early. No question is too basic.

When something goes wrong:

1. Stop using the output.
2. Inform a manager and the AI Governance Lead.
3. Document what happened and what was learned.
4. Update guidance so it does not repeat.
5. If harm occurred, notify leadership for incident response coordination.

Common issues to watch for:

Last updated: February 23, 2026 — [Toby Nwabuogor, AIGP](#) & [Dr. Julie Boone](#) (CC BY-SA 4.0)



- Accepting AI output without verification
 - Using free tools with sensitive data
 - Creating fear instead of a learning culture
-

Policy review

Review this policy **quarterly during the first year**, then **every six months thereafter**.

Set your review dates:

- **Quarterly review:** *Date 1, Date 2, Date 3, Date 4*
 - **Six-month review:** *Date 5...*
-

Approved AI tools

Current approved tools: *[To be completed by organization]*

Free or personal tools

- *[List approved free tools with restrictions, for example: “ChatGPT free tier: Green data only”]*

Enterprise tools

- *[List approved enterprise tools, for example: “Microsoft Copilot for Business: Green and Yellow data”]*

To request a new tool, contact *[AI Governance Lead]* with: tool name, purpose, data type (traffic light color), privacy policy link, and cost.

Contact and questions

AI Governance Lead: *[Name, email, phone]*

When to reach out



- Unsure whether AI is appropriate for a task
- Want to try a new tool or use case
- Something went wrong or needs reporting
- Have ideas for improving this policy
- Need training or support

For toolkit consultation and training: CivicPlay.ai

Acknowledgment

I have read and understand this Responsible AI Use Policy. I agree to follow these guidelines and ask questions when I am uncertain.

Name: ____

Signature: ____

Date: ____



Implementation guide

First 48 hours

Three must-dos:

1. **Customize:** Fill in [Organization Name], [AI Governance Lead], and mission alignment questions.
2. **Brief the team:** Hold a staff meeting, answer questions, and collect acknowledgments.
3. **Set up support:** Approve initial tools and create a question channel (email, Slack, shared doc, and similar).

Build support (as time allows):

- Create a simple FAQ with common questions.
- Document 3 to 5 example scenarios with yes or no reasoning.
- Share 2 to 3 useful prompts.
- Schedule quick check-ins: “What worked?” and “What did not?”

First 90 days

Weeks 1 to 4:

- Focus on low-risk, high-value use cases.
- Document successes and challenges.

Weeks 5 to 8:

- Expand to more use cases.
- Refine guidance based on what you learn.

Weeks 9 to 12:

- Evaluate what is working and what needs adjustment.
- Update the policy if needed.

Success indicators

You will know this policy is working when:

- Staff ask questions without fear



- AI use increases in approved areas
 - Learnings are shared openly
 - Verification becomes routine
 - Innovation happens within guardrails
 - Community trust is maintained
-

Decision rules (fast, consistent answers)

The 60-second decision tree

1. What data is involved? (Green, Yellow, Red)
2. What is the task? Drafting, summarizing, deciding about a person, or communicating externally?
3. Can you verify the output? If no, do not use AI.
4. Who approves? If the work affects others, the approver must be clear.
5. Would stakeholders trust this if they knew? If no, stop and escalate.

Non-negotiables

- No AI tool makes final decisions about people.
 - No personal or restricted data goes into AI tools without an explicit, reviewed exception.
 - All outputs get human review before use.
 - If you cannot verify it, you cannot use it.
-

Verification (how to reduce harm)

General verification checklist

- Accuracy: facts, numbers, names, dates
- Completeness: nothing important missing
- Tone: matches your organization and audience
- Bias: no stereotypes, exclusion, or unfair framing
- Privacy: no identifying details
- Accessibility: plain language and clear structure



Verification triggers (always double-check)

- Anything that sounds overly confident
 - Statistics or legal claims
 - Policy statements
 - Anything involving youth, safety, eligibility, or finances
-

Incident and near-miss handling (no-blame learning)

What counts as an incident

- Sensitive data was shared
- Harmful, biased, or unsafe output
- Incorrect information reached stakeholders
- Unauthorized tool use

What to do immediately

1. Stop using the output.
 2. Notify the AI Governance Lead (and leadership if serious).
 3. Record what happened, how it was caught, and what to change.
 4. Update guidance so it does not repeat.
-

Common challenges and solutions

Challenge: “We’re too small to need a policy”

Reality check: If even one person is using AI with participant or client data, basic guardrails matter.

Solution: This policy takes a few hours to implement. The risk of not having it is much higher than the effort.

Challenge: “I don’t have time to be the Governance Lead”



Solution:

- This role is 2 to 4 hours per month.
- Most of that time is just being available for questions.
- It can be added to an existing role (often ED, admin coordinator, or the most tech-comfortable staff member).
- Alternatively, rotate the role every 6 to 12 months.

Challenge: “Our staff won’t follow this”

Solution:

- Frame it as support, not restriction.
- Start with what people are already doing and make it safer.
- Celebrate people who ask questions.
- Keep it simple. The 4-Question Test is designed to be easy to remember.

Challenge: “What if someone breaks the policy?”

Solution:

- First time: assume good intent and educate.
- Pattern: have a conversation about why.
- Serious breach: treat like any other policy violation.
- Focus on learning, not punishment.

Challenge: “AI is changing too fast for a policy”

Solution:

- This policy is principles-based, not tool-specific.
- The 3Cs and 4-Question Test work regardless of new AI capabilities.
- Quarterly reviews in year one keep the policy current while you learn fast.
- A living document approach allows quick updates when needed.

Challenge: “We can’t afford enterprise AI tools”

Solution:

- Free tools can be fine for Green data.



- Many small orgs operate entirely on free tiers.
 - Upgrade to enterprise only if you need to use AI with Yellow data.
-

Sector scenarios (small org, recreation and community services)

Use these as training prompts and decision practice.

Scenario A: Youth program newsletter draft

- Data: program schedule (usually Green or Yellow depending on whether it is published)
- AI use: drafting is usually fine with review
- Verification: dates, times, costs, accessibility details

Scenario B: Subsidy or financial assistance communication

- Data: likely Red
- AI use: do not paste identifying details
- Safer approach: draft a generic template without personal information, then fill in details manually

Scenario C: Incident report or safety event summary

- Data: often Red
 - AI use: avoid
 - If summarization is needed, consult the Governance Lead for a safe workflow
-

Legal, privacy, and accessibility (plain-language navigator)

When to get help

- Youth or vulnerable population data is involved



- Financial hardship or subsidy eligibility
- Health or medical information
- Automated decisions, scoring, or recommendations about people
- Vendor contracts or unclear data storage or retention

Accessibility basics for AI-assisted communications

- Use plain language.
 - Use clear headings.
 - Avoid jargon.
 - Provide alt text and captions when applicable.
-

Maintenance rhythm (keep it sustainable)

Monthly (2 to 4 hours)

- Review questions asked.
- Review incidents and near-misses.
- Update the approved tools list.

Scheduled reviews

- Review on the cadence in the policy.
 - Also review whenever there is a significant tool or risk change.
-



Policy implementation checklist

Policy foundation

- Policy customized with organizational details
- Mission alignment questions answered
- Governance Lead appointed
- Leadership approval obtained
- All staff briefed and acknowledged

Tools and access

- Approved tools list created
- Enterprise vs. free tools clarified
- Staff know which tools can be used for which data colors
- Access provisioned for approved tools (as applicable)

Support systems

- Question channel established
- Incident reporting process defined
- Governance Lead contact information shared
- Basic FAQ or guidance document created

Learning culture

- Safe reporting culture communicated
- First successes shared
- Learning check-ins scheduled
- Prompt and example library started

Ongoing review

- Quarterly review dates set
- Six-month review date set
- Tracking system in place (simple log or doc)
- Feedback collection method defined

Disclaimer and disclosure

Transparency statement

In the spirit of transparency, the authors disclose the following.

Purpose and intended use

- Provide a practical, principles-based framework for responsible AI use in small organizations.
- Support consistent decision-making, appropriate safeguards, and ongoing learning.
- Balance innovation with the protection of trust, privacy, and safety.

Scope and limitations

- This template is designed for organizations *using* AI-enabled tools, not for developing or deploying AI systems.
- It does not eliminate the need for professional judgment, verification, or context-specific decision-making.
- It cannot prevent all AI-related risks, nor guarantee any specific outcome.
- It reflects the Canadian context and the state of common AI tooling and guidance as of February 2026.
- It assumes basic digital infrastructure and routine organizational controls (for example, appropriate access management and document handling practices).

Policy disclaimer

This document is provided for informational and template purposes only. It is intended as a practical framework and thinking tool, and is not a legally definitive instrument.

No legal advice

- This document does not constitute legal advice.
- Organizations should obtain appropriate legal, privacy, and risk management review prior to adoption.

Implementation responsibility



- Organizations are responsible for adapting this template to their specific context, including governance structure, operational realities, community needs, and applicable legal and contractual obligations.
- AI-related laws, regulations, and guidance are evolving. Organizations should monitor relevant updates and revise their policy and practices accordingly.

Compliance reminder

Organizations adopting this template are responsible for complying with all applicable requirements, which may include:

- Privacy legislation (PIPEDA and relevant provincial privacy statutes)
- Accessibility legislation (Accessible Canada Act and relevant provincial standards)
- Human rights legislation
- Sector-specific regulations and standards
- Employment standards, collective agreements, and internal HR policies

License & Attribution

This toolkit is provided as a resource for Canadian recreation and community service organizations. It is licensed as [CC BY-SA 4.0](#).

You may:

- Use and customize for your organization
- Share with peer organizations (especially other small orgs!)
- Adapt completely to your needs
- Simplify even further if needed

Please:

- Credit authors Toby Nwabuogor and Dr. Julie Booke when sharing
- Tell us how it's working (we want to learn too!)
- Connect with other small orgs & teams using this

Authorship and development disclosure

This policy template was developed through a collaboration between:

- [Toby Nwabuogor, AIGP](#) ([CivicPlay.ai](#))
- [Dr. Julie Booke](#) (Mount Royal University)



Development process:

- Based on ongoing field research with Canadian recreation, sports, and community service organizations conducted by [Dr. Julie Booke](#)
- Informed by the *State of AI Readiness in Canada's Community Services* report ([CivicPlay.ai](#), 2025)
- Informed by *The State of Artificial Intelligence Adoption in Canadian Nonprofits* (Imagine Canada and CCNDR, 2026)
- Aligns with the NIST AI Risk Management Framework and the OECD AI Principles
- Informed by SIRC (Sport Information Resource Centre) articles
- Synthesizes best practices from multiple frameworks and sources including:
 - National AI Centre Australia AI Policy Guide
 - Privacy Commissioner of Canada AI guidance
 - James Kavanagh's AI governance research
 - Sector-specific implementation experience

AI assistance in development: AI tools assisted in drafting, formatting, and synthesizing research. All content was reviewed, verified, and approved by human experts.

Feedback and updates:

We welcome [feedback](#) to improve future versions:

- [Toby Nwabuogor, AIGP](#): toby@civicplay.ai
- [Dr. Julie Booke](#): jbooke@mtroyal.ca

Version history:

- Version 1.0 (February 2026): Initial release

Appendix: related resources (optional)

Consider developing these resources as your AI adoption progresses, but only if they offer tangible benefits to your team.

- **3Cs + 4-Question Test (quick reminder):**
 - **Care:** Protect people and privacy.
 - **Clarity:** Be able to explain and document AI use.
 - **Control:** Humans stay accountable.
 - If you cannot verify it, if it is confidential, if you cannot approve it, or if stakeholders would not trust it, stop and ask the Governance Lead.
- **Traffic Light Data Classification poster:** a visual reminder for staff desks
 - **Green** is public and okay to use, **Yellow** is internal and only okay in approved enterprise tools, and **Red** is sensitive and should not be put into AI tools. When unsure, treat it as Red.
- **Prompt and use case library:** a shared collection of useful prompts and use cases
 - Use case name
 - Inputs (data color)
 - Outputs and audience
 - Verification steps
 - Approval needed
- **Tool approval request form:** a simple form for requesting new tools
 - Tool name and provider
 - What problem it solves
 - What data colors it touches
 - Where data is stored (if known)
 - Decision: approved, approved with conditions, or not approved
- **Near-miss record:** a simple log to capture learning opportunities
 - What almost happened
 - How it was caught
 - What we changed

This is part of the [Responsible AI Governance Toolkit for Canadian Recreation and Community Service Organizations](#), jointly created by [Toby Nwabuogor, AIGP \(CivicPlay.ai\)](#) and [Dr. Julie Booke \(Mount Royal University\)](#).