

Critical Integrity Report

Nonprofit AI Policy Framework - Practitioner Tool Audit

2026-04-17 · policy · hybrid · MCP-Enhanced

CRITICAL INTEGRITY SNAPSHOT

Moderate Integrity

71/100 points



Knowledge Tradition: hybrid

Scoring note: Scored as a practitioner advocacy framework, not a research document. Grounding and framing interpreted relative to the practitioner-tool standard the author intends.

PLAIN-LANGUAGE READ

This framework works well as a practitioner tool for audiences who already sense the problem - but it underdelivers on urgency creation, AI-vs-data-governance differentiation, and standalone actionability for cold audiences.

BEFORE USING THIS TEXT

Before using this text, ask whether your nonprofit audience already feels the AI governance problem - or needs to be convinced it exists and is distinct from what they already do.

Quick-Scan Findings

Top Biases Identified

BIAS NAME	WHERE	WHY IT MATTERS	LENS
Overgeneralization <small>ct:biases:methodology:overgeneralization</small>	AI has fundamentally changed how decisions get made	Presented as settled fact across all nonprofit contexts - but many small orgs have not yet experienced this shift and will not feel urgency from a claim that...	Critical
WEIRD Bias <small>ct:biases:methodology:weird-bias</small>	Let AI support, not replace, the relationships that define nonprofit impact	The model of 'human connection' and 'community trust' reflects Western, educated, institutional organizational contexts. Indigenous-led orgs and newcomer-ser...	Compassionate
Selection Bias	Seven Failure Patterns section	The failure patterns reflect orgs that engaged a consultant - self-selected, with some capacity to seek advisory services. Orgs that govern well without cons...	Critical
Social Desirability Bias <small>ct:biases:methodology:social-desirability-bias</small>	That's not a compliance problem	Framing compliance-first governance as the wrong approach may cause readers to perform alignment with the change-leadership frame without actually adopting i...	Critical

KEY TENSION

The framework is designed for audiences who already feel the AI governance problem - but its stated purpose is to create urgency and action in organizations that don't yet feel it, and those two audiences need differe...

RECOMMENDED NEXT STEP

Add one focused paragraph to the Central Argument section that explicitly names how AI governance differs from existing PIPEDA/data governance compliance - and why basic data go...

Improvement Prompts

To strengthen Grounding & Scope Fit: Add one sentence naming the evidence type - e.g. 'This framework draws on practitioner observation across Canadian nonprofit consulting engagements' - to set appropriate expecta...

To strengthen Assumptions & Context: Add a single paragraph - ideally in the Central Argument - that names the specific ways AI governance differs from PIPEDA/data governance compliance. This is the gap that preven...

To strengthen Bias Transparency: Own the standpoint explicitly. A brief author bio or positionality statement at the start - 'This framework reflects X years of consulting with Canadian nonprofits' - would stre...

To strengthen Framing & Audience Openness: Add 'diverse' or 'inclusive' to the Human Judgment principle language. Add one sentence in Getting Started for micro-orgs (under 5 staff) that names what governance looks like b...

To strengthen Positionality & Power: Add a brief author bio or 'About this framework' note at the start. Consider one sentence acknowledging that Indigenous data governance frameworks (OCAP) and diverse organizatio...

Dimension Deep-Dive

Grounding & Scope Fit 18/20 pts

Well-grounded for its stated scope - Canadian nonprofits, practitioner experience. Broad claims like 'AI has fundamentally changed how decisions get made' outrun the evidence base but are appropriate for a framework making an argument, not a research synthesis.

Finding traced to: Central argument asserts category shift without empirical support, Five Shifts section is the strongest grounding - observational but coherent

To strengthen Grounding & Scope Fit: Add one sentence naming the evidence type - e.g. 'This framework draws on practitioner observation across Canadian nonprofit consulting engagements' - to set appropriate expecta...

Assumptions & Context 13/20 pts

Universal language ('most staff don't understand,' 'AI literacy is becoming professionally essential') does structural work without acknowledgment. The framework assumes the reader already agrees AI governance is distinct from data governance - but that assumption is precisely what skeptical nonprofits will contest.

Finding traced to: 'Most staff don't understand where data goes' - stated as fact, evidence unstated, No explicit case made for why existing data governance is insufficient

To strengthen Assumptions & Context: Add a single paragraph - ideally in the Central Argument - that names the specific ways AI governance differs from PIPEDA/data governance compliance. This is the gap that preven...

Bias Transparency 11/20 pts

The framework presents advocacy as neutral analysis. The 'What Makes This Different' section positions four alternative models as straw alternatives without engaging their strongest versions. The author's consulting standpoint - and its inherent selection bias toward orgs that hired a consultant - is not disclosed.

Finding traced to: Four comparisons position alternatives uncharitably, No disclosure of practitioner standpoint or evidence source

To strengthen Bias Transparency: Own the standpoint explicitly. A brief author bio or positionality statement at the start - 'This framework reflects X years of consulting with Canadian nonprofits' - would stre...

Framing & Audience Openness 16/20 pts

Intentionally concise framing is appropriate for the document's purpose. Small targeted additions - the word 'diverse' in human judgment language, a brief acknowledgment that org size affects governance design - would open the door to broader applicability without bloating the document.

Finding traced to: Alternative interpretations not visible - appropriate given format, addressable with targeted words, No capacity-tier guidance for orgs too small for two-layer governance

To strengthen Framing & Audience Openness: Add 'diverse' or 'inclusive' to the Human Judgment principle language. Add one sentence in Getting Started for micro-orgs (under 5 staff) that names what governance looks like b...

Positionality & Power 13/20 pts

The framework speaks about affected communities but not with them. Service recipients, Indigenous data sovereignty frameworks, and racialized or disabled staff are absent from both the evidence base and the governance model. The author's consulting standpoint is implied but unnamed.

Finding traced to: Pillar 3 'human connection' model reflects Western professional org assumptions, No positionality statement or author background disclosed

To strengthen Positionality & Power: Add a brief author bio or 'About this framework' note at the start. Consider one sentence acknowledging that Indigenous data governance frameworks (OCAP) and diverse organizatio...

Assumption Audit

ASSUMPTION	TESTABLE?	WHAT WOULD TEST IT
The reader already agrees that AI governance is distinct from basic data governance compliance	testable	A nonprofit that says 'our data governance covers this' would test it - and the framework currently has no answer for them
Named accountability ownership improves governance outcomes regardless of org size	testable	Comparison of governance outcomes in micro-orgs with named vs. distributed accountability
The Seven Failure Patterns represent the full distribution of nonprofit AI governance approaches	testable	Evidence from orgs that self-governed successfully without consulting engagement
AI literacy becoming professionally essential is a settled fact rather than an argument	not testable	This is the document's core claim - it is the argument, not a premise requiring support

Alternative Frames

Cynefin Framework - Complex vs. Complicated

ct:thinking-lenses:systems:cynefin-framework

The framework treats AI governance as a complicated problem (expert knowledge + best practices = right answer). Many nonprofits are in the complex domain (emergent, context-dependent). This reframe suggests the Eight Steps may need a probe-sense-respond variant for orgs in early chaos.

Three Cs lens: Critical

Cross-Cultural Anti-Flattening

ct:methodology-checks:preflight:cross-cultural-anti-flattening

The framework's universalizing language flattens genuine organizational diversity across Canadian nonprofits. Indigenous-led orgs, Francophone organizations, and newcomer-serving agencies have distinct accountability and relational models that the current framing doesn't accommodate.

Three Cs lens: Compassionate

Argument Map

Central claim: AI policy is a change leadership problem disguised as a technology problem - and Canadian nonprofits need a mission-centered, capacity-aware governance framework rather than compliance checklists.

Claim type: policy · Confidence: medium

The framework argues that AI policy fails when treated as a compliance or technology problem, and succeeds when treated as a change leadership problem. It provides a structured practitioner system for Canadian nonprofits to build governance that protects mission, community trust, and charitable status.

Supporting Reasons

- AI has changed how decisions, knowledge, trust, and risk work inside organizations (Five Shifts)
- Seven documented failure patterns show that compliance-first approaches don't work
- Canadian legal context (no AI-specific law) creates both freedom and responsibility to self-govern
- The digital divide inside organizations makes AI literacy a workplace equity issue

Objections

- Basic data governance (PIPEDA) may already cover what nonprofits need - the framework doesn't directly rebut this
- Two-layer governance assumes organizational capacity many small nonprofits don't have
- The Getting Started section is thin for standalone actionability without the paid follow-up

Missing Voices

- Service recipients and community members who receive nonprofit services
- Indigenous organizations operating under OCAP and data sovereignty frameworks
- Micro-orgs (under 5 staff) who cannot implement the two-layer model

Evidence Gaps

- No concrete case made for how AI governance differs from existing data governance compliance
- No evidence of outcomes from organizations that adopted this framework
- Failure patterns drawn from consulting-engaged orgs - selection bias unacknowledged

Argdown Export

(AI policy is a change leadership problem disguised as a technology problem - and Canadian nonprofits need a mission-centered, capacity-aware governance framework rather than compliance checklists.)

- + AI has changed how decisions, knowledge, trust, and risk work inside organizations (Five Shifts)
 - + Seven documented failure patterns show that compliance-first approaches don't work
 - + Canadian legal context (no AI-specific law) creates both freedom and responsibility to self-govern
 - + The digital divide inside organizations makes AI literacy a workplace equity issue
 - Basic data governance (PIPEDA) may already cover what nonprofits need - the framework doesn't directly rebut this
 - Two-layer governance assumes organizational capacity many small nonprofits don't have
 - The Getting Started section is thin for standalone actionability without the paid follow-up
 - ? The reader already accepts that AI governance is distinct from data governance
 - ? Named accountability ownership works across all org sizes
 - ? The Seven Failure Patterns represent the full distribution of governance approaches
 - ! No concrete case made for how AI governance differs from existing data governance compliance
 - ! No evidence of outcomes from organizations that adopted this framework
 - ! Failure patterns drawn from consulting-engaged orgs - selection bias unacknowledged

Evidence Needed Next

Check-worthy means worth verifying, not false. Only factual claims receive verification priority.

CLAIM	TYPE	PRIORITY	EVIDENCE NEEDED
<p>Most staff don't understand where data goes when they use AI tools</p> <p>Principle 3 - Data Governance Must Be Explicit</p>	factual	medium	<p>Survey data or sector research on staff AI literacy across Canadian nonprofits</p> <p>Likely sources: sector surveys; work-force research; nonprofit technology reports</p> <p>Suggested queries: Canadian nonprofit staff AI literacy survey; nonprofit employee data governance awareness</p> <p>Support/contrast: needs_evidence</p> <p>Uncertainty: Likely directionally correct but scope is overstated - some orgs have sophisticated data governance training</p>
<p>In the nonprofit sector, trust erosion is an existential threat</p> <p>Canadian Legal & Regulatory Context section</p>	causal	not_ranked	<p>Cases of Canadian nonprofits that experienced AI-related trust erosion and their outcomes</p> <p>Likely sources: sector news; case studies; charity sector reports</p> <p>Suggested queries: nonprofit trust crisis AI Canada; Canadian nonprofit AI reputational harm</p> <p>Support/contrast: not_ranked_for_truth</p> <p>Uncertainty: True for the most vulnerable org types this framework serves; overstated as universal</p>
<p>AI policy is a change leadership problem disguised as a technology problem</p> <p>Central Argument section</p>	normative	not_ranked	<p>Comparative outcomes from orgs using change-leadership vs. compliance-first AI governance approaches</p> <p>Likely sources: implementation records; sector research; case studies</p> <p>Suggested queries: change management AI governance nonprofit outcomes; compliance vs culture AI policy effectiveness</p> <p>Support/contrast: not_ranked_for_truth</p> <p>Uncertainty: This is the document's core argument - it is a values claim, not</p>

CLAIM	TYPE	PRIORITY	EVIDENCE NEEDED
			a factual one, and should be evaluated as such

Questions To Bring To This Text

Questions to bring to your next conversation about this text:

- Could you hand this to a nonprofit executive director who currently thinks 'our data governance covers AI' and have them walk away convinced otherwise - and if not, what's missing?
- What does governance look like for a 3-person nonprofit that cannot implement two-layer structure - and does the framework leave them better or worse off than before they read it?
- If the paid follow-up is gated behind a paywall, does this document give enough for a nonprofit to take one concrete action - or does it primarily create awareness of a problem without a path forward?
- What specific scenario or example would make the AI-vs-data-governance distinction immediately clear to a skeptical nonprofit board?
- How does the framework handle orgs that have already adopted AI tools widely - is 'Getting Started' the right frame for them, or do they need a different entry point?
- What would Sarah say to a nonprofit that reads this and concludes 'this is good but we'll wait for the paid guide before doing anything' - is that an acceptable outcome for this document?

Next Steps For Deeper Thinking

1. Add a focused paragraph to the Central Argument that names at least two concrete ways AI governance differs from PIPEDA/data governance compliance - this dir...
2. Add a brief positionality or author bio note at the start of the document - one paragraph naming Sarah's background and where this framework comes from would...
3. Add capacity-tier guidance for micro-orgs (under 5 staff or no formal governance structure) in the Getting Started section - even two sentences naming what g...

Reasoning Trap Checks

Causation-leap risk: Risk of reading timing or association as if it proves direction or cause.

Next check: Ask what comparison would separate coincidence from cause.

Missing-denominator risk: Risk of treating a number as meaningful before knowing the denominator, sample size, or comparison group.

Next check: Ask what the number is out of and which comparison group makes it interpretable.

Comfortable-answer risk: Risk of treating the easiest answer as settled before pressure-testing alternatives.

Next check: Ask what evidence would make this answer harder to defend.

Binary-frame risk: Risk of collapsing a messy decision into a yes-or-no frame and missing a middle path.

Next check: Ask what third option, delay, or partial step has been left out.

Base-rate neglect risk: Risk of judging the claim without comparing it to the normal rate or background pattern.

Next check: Ask what baseline rate would make this finding look large, small, or ordinary.