


WHITEPAPER

AI Governance for Public Sector Project Leaders

Principles, Frameworks, and Ontario-Specific Guidance for Responsible AI in the Public Sector PMO



Avishek (Avi) Chaudhury

MBA · PMP · PROSCI Certified Change Practitioner · CSM · SAFe Agilist

Strategic Advisor, Enterprise PMO Governance & Portfolio Decision Intelligence

Executive Summary

A strategic briefing for public sector project managers, PMO directors, and digital transformation leaders navigating the responsible deployment of artificial intelligence.

I want to start with an honest admission: for much of my career, the phrase "AI governance" would have made my eyes glaze over. It sounded like the kind of thing that lived in a policy document no one read — produced by a committee that met quarterly and accomplished roughly nothing between meetings.

Then I watched a well-intentioned AI tool in a public sector environment quietly produce biased recommendations for three months before anyone noticed. The tool was doing exactly what it was designed to do. The problem was that no one had defined what it was supposed to be accountable for. That's a governance failure — and it's more common than any of us want to admit.

Across Ontario's public sector, AI tools are being embedded into project delivery, procurement risk analysis, citizen service, and reporting workflows. The productivity case is real. But productivity without accountability is just faster failure. And in public service, faster failure has real consequences for real people.

This whitepaper argues that AI governance is not an IT function to be handed off and forgotten. It is a project management discipline — one that belongs squarely in the hands of the people who are accountable for delivering outcomes. That's us.

Drawing on Ontario's emerging regulatory environment — including EDSTA, the IPC's AI guidelines, and OHRC principles on algorithmic fairness — this paper presents a practical, five-pillar governance framework that PMO leaders can actually use.

Not a compliance checklist — but more of a leadership tool.

CORE ARGUMENT

Across 14 OECD nations, fewer than 1 in 3 public sector project managers feel confident assessing AI risk in their projects. Ontario's regulatory environment is tightening.

The window to build this competency — before it becomes mandatory rather than strategic — is now.

1. The AI Inflection Point in Public Sector Project Delivery

Why AI governance is no longer optional — and why the people in the room must own it.

1.1 From Pilot to Plumbing

For most of the past decade, AI in the public sector meant experimental chatbots and isolated analytics proofs-of-concept that got demoed at conferences, praised in press releases, and quietly retired six months later. That era is ending.

Tools like Microsoft Copilot are already embedded in the SharePoint and Teams environments that Ontario municipal and provincial teams use daily. AI is no longer a pilot — it has become operational plumbing. Which means governance can no longer be a pilot either.

I sometimes put it this way at workshops: If you wouldn't deploy a new procurement system without a control framework, why would you deploy an AI system that influences procurement recommendations without one?

The answer is usually uncomfortable silence. Which is exactly the right answer.

1.2 The Accountability Gap

Here is the structural problem: organizations are deploying AI capabilities — often driven by vendor contracts or executive enthusiasm — faster than they are establishing accountability frameworks to oversee them. The result is an accountability gap: AI systems are influencing consequential decisions, but no one person or team owns the outcomes those decisions produce.

In the public sector, this gap is especially dangerous. Unlike private organizations that can absorb a governance failure as a reputational cost, public institutions are bound by procedural fairness, transparency, and equity obligations under MFIPPA, the Ontario Human Rights Code, and increasingly, EDSTA.

An AI system that influences resource allocation, project prioritization, or service delivery is subject to those obligations whether anyone remembered to document that it is or not.

HEADS UP


One of my professors had a phrase that stuck with me: 'Accountability doesn't transfer with delegation.' It applies perfectly here.

*You can outsource the AI capability to a vendor. **You cannot outsource the accountability for what it does.***

1.3 What the Data Is Telling Us

The numbers are not flattering — but they are clarifying.

Metric	Finding	Source
Public sector AI adoption (Canada)	61% using AI in operations	KPMG Canada, 2025
Orgs with formal AI governance policy	28% of public sector	Deloitte Public Sector, 2025
PM confidence in AI risk assessment	31% report high confidence	OECD, 2024
Projects with AI-specific risk registers	17% of surveyed PMOs	PMI Pulse, 2025



The pattern here is consistent: AI adoption is running at roughly twice the rate of governance readiness. As someone who has spent 15 years in environments where the gap between capability and accountability creates real harm — from transit delivery failures to healthcare systems under pressure — I find these numbers clarifying rather than alarming. We know where the problem is. Now we can do something about it.

2. Ontario's Regulatory Landscape for AI

The rules of the road — and why project managers need to know them, not just lawyers.

2.1 Enhancing Digital Security and Trust Act (EDSTA)

Ontario's Enhancing Digital Security and Trust Act, passed in 2024, establishes the foundational legislative framework for how the Province and its broader public sector must approach digital systems — AI included. For PMO leaders, the practical implications are:

- Algorithmic systems used in government decision-making processes may require disclosure and documentation.
- Risk assessments for high-impact digital systems — including AI tools embedded in project management platforms — may be required before deployment.
- Data handling obligations under EDSTA intersect with existing MFIPPA requirements, creating dual compliance considerations that most project teams have not yet mapped.


My working rule: if you wouldn't be comfortable explaining how this AI tool works to an auditor, a journalist, or your project sponsor — you haven't done enough governance work yet.

2.2 IPC Artificial Intelligence Guidelines

The IPC has published guidance on AI in public sector contexts with emphasis on privacy by design and algorithmic accountability. The three principles I return to most often in practice:

- AI systems processing personal information must be assessed for privacy impact before deployment. Not after. Before.
- Automated decision-making must be transparent: individuals affected have a right to understand how decisions affecting them were made.
- Data minimization applies: AI systems should not consume more personal data than strictly necessary for their stated purpose.

In my healthcare years, we called this principle **'minimum necessary access.'** The same logic applies to AI systems — they should touch only what they genuinely need to do their job.



2.3 OHRC Principles on Algorithmic Fairness

The Ontario Human Rights Commission has made clear that the Ontario Human Rights Code applies to algorithmic systems. Where AI tools influence decisions affecting protected groups — resource allocation, vendor selection, workforce planning — organizations can be held liable for discriminatory outcomes even when those outcomes were produced by an algorithm rather than a human decision-maker.

This is not a theoretical risk. It is a live accountability question on every project where AI tools are being used to influence how people, resources, or opportunities are distributed. As someone who founded an organization dedicated to inclusion for people on the autism spectrum, I hold this principle personally as well as professionally.

Equity is not a compliance checkbox - it is a design requirement.

REGULATORY WATCH

Ontario's AI regulatory landscape is actively evolving through 2026-2027. PMO leaders should build a standing mechanism to monitor regulatory updates — not treat this as a one-time read.

Think of it like your risk register: something you return to regularly, not something you file and forget.

3. The Five Pillars of Public Sector AI Governance

Turning principles into practice — a framework grounded in 15 years of public sector delivery.



Before I present this framework, a disclaimer that I mean genuinely: frameworks are only useful if they change behaviour. I have seen enough governance documents produced and promptly forgotten to know that a five-pillar model is only worth the paper it's printed on if it shows up in your project charters, your risk registers, your stage gate reviews, and your conversations with sponsors.

These five pillars are not abstract. Each one maps directly to something you already do — and asks you to do it with AI in mind.

01

Accountability — Every AI Action Needs a Human Owner

Every AI-influenced decision in a project must have a named human accountable for its outcomes. This is non-negotiable in the public sector. AI tools should appear in your RACI, with clear delineation between where AI informs and where humans decide. If you cannot name the person accountable for an AI system's outputs, you are not ready to deploy it.

02

Transparency — If You Can't Explain It, You Can't Govern It

Project teams must be able to explain, in plain language, how AI tools are being used, what data they consume, and what decisions they influence. I use a simple test: can I explain this AI system's role in our project to my project sponsor in under two minutes? If not, something is wrong — either with the system or with our understanding of it. Explainability isn't just a regulatory obligation. It's a project health indicator.

03

Fairness — Equity Is a Design Requirement, Not an Afterthought

AI tools used in public sector projects must be assessed for potential disparate impact on equity-deserving groups, consistent with OHRC guidance. This requires project managers to include fairness criteria in vendor assessment scorecards, procurement specifications, and acceptance testing criteria. The Living Autism Foundation taught me something that applies here: systems designed without equity in mind rarely become equitable through retrofit. Build it in from the start.

04

Safety — Risk-Proportionate Oversight for Every AI System

AI systems should be governed proportionate to their risk level. An AI-assisted scheduling tool needs lighter oversight than an AI system influencing procurement decisions. PMOs should adopt a tiered AI risk classification aligned to their enterprise risk framework. My rule of thumb from healthcare: if the AI's recommendation, if wrong, could harm a person — that's a high-risk system. Govern it accordingly.

05

Continuous Oversight — AI Governance Is a Verb, Not a Noun

AI systems change behaviour as they process new data. A governance framework that was appropriate at deployment may not be adequate six months later. Governance must be embedded into ongoing project monitoring cycles — not a one-time assessment. I think of it like my Tae Kwon Do training: you don't earn a belt and then stop practicing. You keep showing up. AI governance is the same.

4. Embedding AI Governance into the PMO Operating Model

This is where frameworks become discipline.



4.1 Initiation & Charter Phase

AI governance begins at project inception — not at the post-implementation review when things have already gone sideways. Before an AI tool is selected or procured, the project charter should address:

- AI Tool Registry Entry: What AI tools are anticipated? What data will they access? Who reviewed and approved each one?
- Accountability Assignment: Who is the named AI Accountability Owner for this project? Not a committee. A person.
- Regulatory Touchpoint Analysis: Does this project trigger EDSTA, MFIPPA, or OHRC considerations?
- Equity Impact Screening: Does this project involve decisions that could affect equity-deserving groups?

At the City of Toronto's Transit Expansion Office, we learned early that questions asked at the charter stage cost almost nothing. The same questions asked after go-live cost considerably more — in time, in trust, and sometimes in outcomes.

4.2 Planning Phase

During planning, governance requirements should become concrete project artifacts. Not policies. Artifacts — things that live in your project file and get reviewed at your next stage gate:

- AI Risk Register: A dedicated AI risk section, separate from or integrated into the project risk register. This is not optional for execution-tier AI systems.
- Vendor AI Governance Assessment: Procurement criteria should include AI transparency, data governance, and bias testing requirements. If your vendor cannot answer these questions, that is a procurement risk, not just a technical concern.
- Stakeholder AI Communication Plan: How will project stakeholders be informed about AI use? In public sector contexts, this sometimes includes the public.

4.3 Execution & Monitoring Phase

During execution, AI governance becomes an operational discipline embedded in your regular rhythms:

- AI Performance Dashboard: Track AI system indicators alongside project KPIs. If your project dashboard has twelve metrics and zero of them relate to your AI tool's performance, that is a governance gap.
- Human Override Log: Document instances where human judgment overrode AI recommendations. I cannot stress this enough: an override log is governance evidence, not a failure signal. It proves your humans are in the loop.
- Incident Response Protocol: Establish clear escalation paths before you need them, not while you are in the middle of an incident at 4:45 on a Friday afternoon.

4.4 Closure & Lessons Learned

Project closure should include a dedicated AI retrospective alongside your standard lessons learned:

- Did AI tools perform as expected? Were there accuracy or fairness issues that surfaced during the project?
- What governance gaps were identified? How should the next project's AI governance approach be strengthened?

- What data was consumed? Was it disposed of appropriately under MFIPPA and EDSTA?

✓ BEST PRACTICE

The most effective PMOs treat AI governance not as an add-on workstream, but as a lens applied across existing project management disciplines. You do not need a new methodology.

You need to ask different questions inside the one you already have.

5. Building AI Governance Competency in Your PMO

Developing the capability that frameworks require — starting with you.

5.1 The AI-Literate Project Manager

Let me tell you what AI governance competency does not require. It does not require you to understand transformer architecture. It does not require you to be able to explain backpropagation at a dinner party. It does not require a computer science degree, a data science certification, or a subscription to any particular AI newsletter.

What it does require:

- **AI Literacy:** Understanding how common AI tools work at a conceptual level. What is the difference between a machine learning model, a generative AI tool, and a rules-based automation? If you can explain that distinction in plain language, you are ahead of most people in the room.
- **Risk Sensing:** The ability to recognize when an AI application crosses from low-risk administrative territory into high-risk decision territory — and knowing which governance requirements that triggers.
- **Regulatory Awareness:** Familiarity with Ontario's legislative framework and the judgment to know when you need a lawyer in the conversation.
- **Equity Lens:** The habit of asking, for any AI-influenced decision: who might be harmed or disadvantaged by this? And how would we know if that was happening?

I developed most of my AI literacy not through formal training, but through the discipline of asking uncomfortable questions in rooms where it would have been easier to stay quiet. That's a competency you can build starting today.

5.2 PMO-Level Governance Infrastructure

Individual competency must be supported by PMO-level infrastructure — the kind that survives a team member's departure and doesn't live only in one person's head:

- **AI Governance Policy:** A standing PMO policy establishing minimum requirements for AI tool use across all projects. One page is better than twenty. A policy no one can find is the same as no policy.
- **AI Governance Toolkit:** Standard templates for AI risk registers, vendor assessments, transparency reports, and incident logs.

- **AI Governance Review Process:** Integrated into existing stage gate reviews — not a separate process that competes with everything else already on the agenda.
- **AI Community of Practice:** A cross-project forum for sharing what's working, what isn't, and what the regulatory environment is doing. The field is moving fast enough that shared learning is a strategic asset.

5.3 The Director-Level Imperative

For those targeting Director-level and above roles in public sector digital transformation, healthcare, transit, or municipal government — and I have been tracking this across the roles I have applied to and consulted on over the past two years — AI governance literacy is rapidly becoming a baseline expectation, not a differentiator.

Executives are looking for leaders who can translate AI capability into governance accountability. Who can represent the organization's AI posture to oversight bodies and Boards. Who can build the internal systems and team competencies that make AI deployment sustainable — not just fast.

That space — where project management leadership and digital transformation leadership converge — is where the next generation of public sector Directors will be built. I believe that with conviction. And I am betting my career trajectory on it.

6. A Call to Action for Ontario's Public Sector PMO Community



I want to close with something that isn't in any framework document.

I have two boys on the autism spectrum. Watching them navigate systems — educational systems, healthcare systems, community systems — that were designed without people like them in mind has made me constitutionally incapable of treating governance as a bureaucratic obligation. Governance is what decides whether the systems we build include people or exclude them. Whether they amplify human potential or quietly undermine it.

AI in the public sector will do both, depending on how it's governed. We get to decide which outcome is more likely. That is not a small responsibility. But it is ours.

Action 1: Audit Your Current AI Exposure

Identify every AI tool in use across your project portfolio. Assess each against the five-pillar framework. Identify your highest-risk exposure and address it before the next stage gate. If you don't know what AI tools your team is using, that is the first governance gap to close.

Action 2: Establish One Governance Artifact This Quarter

You do not need a complete AI governance framework before you start. Pick one artifact — an AI risk register template, a vendor AI assessment checklist, a transparency reporting format — and implement it on your next project. Perfect is the enemy of done and done is the enemy of the 28% of public sector organizations that currently have no AI governance policy at all.

Action 3: Bring Your Peers

Share this framework with your peer leaders. Advocate for AI governance to be included in your organization's next PMO maturity assessment. The field learns faster when leaders share what they know rather than protect it. That's true in Tae Kwon Do — the senior belts teach — and it's true here.



Avishek (Avi) Chaudhury has spent more than 15 years building the governance frameworks, executive dashboards, and decision-support structures described in this paper — not theorizing about them.

His work has spanned enterprise PMO governance and strategic initiatives leadership across the City of Toronto's Transit Expansion Office, Metrolinx, the Government of Ontario's Ministry of the Solicitor General, and major Ontario healthcare systems including Trillium Health Partners and Michael Garron Hospital, alongside founding and directing The Living Autism Foundation (Canada), building the organization's board governance and reporting structures from the ground up.

His work has included designing enterprise PMO governance frameworks, stage gates, and decision hygiene standards for large infrastructure and transit portfolios; rebuilding executive reporting architecture to cut reporting cycle times and developing governance guardrails for responsible AI adoption within regulated, politically sensitive environments.

Avi can be reached at **avishekchaudhury@yahoo.com**

© 2026 Avi Chaudhury. This whitepaper is intended for educational and professional development purposes. Nothing herein constitutes legal advice. Readers are encouraged to seek qualified legal counsel for matters related to regulatory compliance.