

The Honourable Evan Solomon

Minister of Artificial Intelligence and Digital Innovation

The Honourable Mélanie Joly

Innovation, Science and Economic Development Canada
Government of Canada

September 23, 2026

Subject: OpenMedia's Inputs to 'Have your say on advancing AI transparency in Canada'

Dear Minister Solomon and Minister Joly,

I am writing on behalf of OpenMedia, a community-based non-profit organization dedicated to keeping the Internet open, affordable, and surveillance-free. Our community of over 230,000 people across Canada is deeply invested in shaping informed, participatory, and rights-respecting digital policies.

We appreciate the Government of Canada's initiative to launch this public consultation to gather expertise and lived experiences directly from more diverse stakeholders. A participatory process of this kind is important to ensure Canada understands the problem correctly before we even try to find an answer.

AI itself is powerful, and growing more so. If Canada permits corporations having to have unchecked power over this tool, it could become very destructive. We've witnessed how foreign Big Tech exploited Canadians's data and privacy alike for their interests in the past decades unassisted by advanced AI. Our relationships with reality are increasingly distorted: misinformation and polarization proliferates with the aid of engagement-maximizing algorithms, and threatens the health of our democracy. New AI technologies are bringing these consequences into the physical world, as automation in the workplace shifts the labour markets for better and worse; and mounting damage to the environment and rising energy costs from data centres risk catastrophic consequences that none of us can escape.

As Canada continues to prioritize innovation in a strange time where our country's digital sovereignty is constantly being challenged, making new technologies accountable to the public is one of the most important pillars we must defend. Transparency is a key component of this pillar, but there are other fundamental elements, including respecting rights, privacy and sustainability that together are needed to prevent this pillar from sinking.

In this letter, we will translate key insights from a survey we conducted with over 3,000 Canadians in August 2025, and share policy recommendations that will help Canada make progress in establishing frameworks that give Canadians the protections we deserve.¹

¹ OpenMedia's Inputs to a Renewed National AI Strategy that Serves Canadians – [OpenMedia](#)
OpenMedia Engagement Network

AI-Generated Content

Eroding public trust driven by vast quantities of online content of uncertain origin and unclear quality is already a serious problem: it will be worsened by AI content, which is rapidly worsening our ability to tell what is real or not anymore. Labelling AI-generated content is one direct though inadequate way to mitigate the problem. While **visible watermarks** help ordinary people better understand how, when, why, and in what ways AI was used in the content, **invisible watermarks or provenance metadata** can help protect creative integrity by catching content that attempt to claim false sourcing or authenticity. Some major platforms and AI companies already implemented these methods — e.g. Claude's text watermarking survives copy-and-pasting, Instagram automatically applies an AI content label, and Substack integrates an AI transparency tool to help creators disclose their use of AI in their content. However, inaccuracy and mislabeling will persistently undermine the effectiveness of such methods. We recommend **standardizing a framework with key identifiable icons and categories across platforms** to help Canadians better differentiate AI-generated and AI-assisted content from human-produced content. As this may prove more useful for misinformation than AI-assisted disinformation campaigns, we further suggest government investment in and promotion of chain of custody efforts for non-AI content such as the [Content Provenance and Authenticity](#) voluntary standard.

Canadians' growing uncertainty about content online, whether it is AI-generated or not, is the starting point of an inability to separate relatively accurate information from falsehood that will ultimately undermine our democracy. We are increasingly exposed to content that mislead, manipulate, and erode our confidence and judgement; this problem cannot be easily solved by slamming labels on content. To more deeply address this root cause, Canadians' **digital literacy needs to be constantly renewed; and support for independent journalism is the best way to achieve this**. We ask the government to expand public subsidy support for a diverse range of local, independent, fact-based media, as well as measures such as expanding tax credits or issuing universal media vouchers to encourage Canadians' subscriptions to accredited outlets of our choice.

Sloppy AI content can be harmless and entertaining sometimes,² and in the spirit of freedom of expression, we certainly do not want to see Canada's government police everything online. But an entirely laissez-faire approach to AI is neither responsible, nor in keeping with traditional Canadian values of responsible risk-based governance, nor does it match the asks of our 2025 AI consultation, which favoured strong risk-based regulation today. Our survey respondents asked Canada to adopt concrete steps to establish proactive, inclusive and accountable legal frameworks that protect Canadians and uphold our values. This includes introducing **ex-ante regulation**, that is precautionary legal measures that forbid AI from engaging in clearly harmful and manipulative behaviours, and implementing a risk-based tiered regulatory system, that applies tighter regulation to AI purposes that have higher social risks.

² Older adults know AI is slop. They just like it – [Rest of World](#)

AI Interaction

Transparency around AI interactions may enable Canadians to make more informed decisions, but this will only matter if meaningful choices are available to exercise once disclosure happens. Sadly, this is not the case with digital services concentrated in only a handful of Big Tech companies in many of Canada's current digital services, and AI is headed down the same path. And in Canada, we have one of the most consolidated telecommunications markets among developed countries, where the Big Three telecom giants own most of the networks and services in the country. AI is developing a similar concentrated dynamic, with models based on Gemini, ChatGPT, or Claude underpinning most deployments popular in our country currently.

Where similar frontrunner monopolistic tendencies exist, and customer switching costs are high, Canada should be looking at similar remedies to those we've deployed imperfectly in other sectors. Disclosure-for-informed-decision-making measures will fall short as a remedy. **What Canadians need in these contexts is not just notice, but control:** the regulated ability to consent, or refuse consent, to how their data is used, to contest a decision an AI system made about them, and to know that a human remains accountable for the outcome. All of these require an auditable, designed feature built in the system from the start, not as an afterthought.

At the same time, the Government should be cautious about defaulting to more disclosure as the answer to every gap. As the discussion paper identified, "disclosure fatigue" would only push Canadians to stop seeing disclosure statements, much as we currently are forced to pretend to accept consent agreements, which makes the act of disclosure meaningless. Therefore, active consumer-facing disclosure should be reserved for moments where actions related to consequential decisions, data use, and points where a human alternative or appeal exists, with consumers otherwise given a regulated right to set default permissions for their data and consent that AI systems must read and respect.

Information about AI Systems

Different audiences hold varying stakes in transparency of AI systems. As noted in the discussion, end-users' interests on information could be drastically different from that of researchers. **To our community, the premise of AI systems transparency is about public accountability.** That includes how Canadians' data is being collected, used and trained in models. We urge the government to **put the public interest** — particularly the protection of Canadians' rights and privacy protections — **at the front and centre of any transparency framework**, as the necessary foundation on which AI's economic benefits can be responsibly realized. All members of our society hold their respective, legitimate reasons that justify why and what they need. When establishing frameworks that enhance AI systems, we urge the government to consider and provide multiple levels of accountability to meet all stakeholders' needs. This can include requiring plain language summaries of AI systems accessible to everyone, technical public summaries available for more detailed public analysis, and full access to otherwise confidential material for accredited independent auditors, at both a regulator and vetted academic institutions.

Considering the rapid growth of corporate power and control over the technology, it is critical for Canada to **establish frameworks that scrutinize the full AI value chain**. Accountability cannot stop at developers and deployers. Executives and investors dictate the strategic directions and shape the outcomes of an AI product in ways that a developer or a deployer cannot. **This inherently makes the potential harms of an AI system harder to address, with responsibility distributed across a chain** of design, funding, and deployment decisions, no single link of which may appear negligent in isolation. Since harm can no longer be traced easily to a single actor after the fact, oversight needs to be continuous and documented at each stage.

Recognizing this emerging reality, Canada needs to establish **frameworks that mandate reporting obligations, audits, and designate an independent regulatory body to monitor disclosure and accountability** across the full value chain. We recommend four mechanisms. First, organizations developing or deploying AI systems in high-risk contexts should be required to designate a senior accountable officer who is responsible for compliance. Second, those organizations should publish a risk management framework, attested to at the board level, describing how they identify and mitigate risks. Third, design and deployment decisions should be documented in records available on request, so oversight does not depend on reconstructing events after harm occurs. Lastly, high-risk systems should be subject to independent third-party audit, with results available to both a regulator and the general public.

Effective rights-driven regulation will only exist with a strong and independent regulator charged with its enforcement, and Canada's current plans will not lead to that body. The Digital Safety and Data Protection Commission proposed by Bill C-36 would sit within a ministry whose principle interest is promoting Canada's economy and encouraging AI investment, an obvious contradiction in interests. We do not believe C-36 should remove the role of the independent Office of the Privacy Commissioner in enforcing privacy rights within AI-related matters. If it nonetheless moves forward, the new body must have guaranteed security of tenure for its leadership once appointed, its own appropriation of funds for its activities not under ISSED oversight, a clarified and simplified set of priorities not muddled by the current section 77, and a direct reporting relationship to Parliament, not any ministry.

AI Incidents

AI products are being widely deployed across sectors. Their adoption has tripled since 2024, as noted in the discussion. The risk associated with AI is commensurately expanding—yet Canada currently has no comprehensive law governing how AI should be used across industries, and none on the current horizon. Conventional safeguards, such as human oversight, internal reporting, and the like, are being left to the individual practices of businesses or industries, which make the impact of an AI incident even harder to estimate or predict in advance. And because AI systems increasingly interact across sectors and platforms, an incident's effects may not stay confined to where it originated.

To mitigate AI incidents of this nature and give Canadians the protections we deserve, the government must establish two complementary mechanisms: **a systemic reporting regime** that Canada currently lacks, and an **ongoing government capacity to study AI risks** and

OpenMedia Engagement Network

1424 Commercial Dr, Box 21674, Vancouver, BC, V5L 5G3, Canada // 1.888.441.2640

publish its findings so Canadians and local businesses can better understand and respond to fast-moving systemic risk with the right strategy. This need is growing more urgent as incidents increasingly involve AI agents or decisions made without adequate – or even any – human oversight.

The mandatory reporting regime should **assign responsibility according to who holds control at each stage** of developing and deploying an AI system. Developers who alone see training, design choices, and known failure modes, should hold primary responsibility for testing and preventing harm. Deployers who substantially modify a model's default behaviour should inherit that same responsibility for what they have changed; they should also be responsible for handling day-to-day harm reporting given their close proximity to the system in operation.

This reporting duty on AI harm incidents should follow a government-mandated format, describing what happened, the type of harm it caused, the population affected, whether the system was operating as intended or failed, and what redress and mitigation is being implemented afterwards. The concept of “real risk of significant harm” in PIPEDA could be adapted for use around AI harms. Reporting requirements should cover high-risk near-miss incidents, with penalties for failure to report, and reduced penalties for harms when incidents are reported in good faith. These reports should be provided to the designated regulator in full, with as much information as possible being provided to the Canadian public while respecting the privacy of individuals in Canada.

The framework should include clear definitions that help stakeholders along the AI value chain to identify where harm stems from, who should be responsible for a flaw in a system, whether it is disclosed, foreseeable, preventable or not; and assist in determining what, where and on whom the penalty should fall. The framework should enable a **direct path for civil society and individuals to report suspected violations and harms to the regulator**, rather than relying solely on self-reporting.

To avoid duplicating costly efforts and diffusing our actual regulatory impact, **Canada should partner with domestic and international institutions** for its study of AI risks, and align our reporting frameworks with emerging international standards in like-minded democratic peers. Canada's close and deepening relationship with the European Union, including its prospective status as the EU's first “associate member,” presents a timely opportunity to build interoperability into any Canadian framework from the outset.

AI Agents

Agentic AI use is proliferating among Canadian organizations. In a recent study conducted by Oxford Economics and SAP,³ 200 Canadian business leaders were surveyed and revealed that 66% of their businesses are already piloting agentic AI, yet a staggering 97% admitted that they are not fully prepared to deploy and govern these powerful tools. The ways in which they set strategy for AI adoption are fragmented — most remain department-led (39%) or piecemeal (35%) approaches, while only a few are taking centralized, strategic approaches (18%). With

³ Canadian Businesses Race to Adopt Agentic AI, but 97% Say They Aren't Fully Ready – [PR Newswire](#)
OpenMedia Engagement Network

their own governance efforts embryonic, and no external disclosure or accountability requirements, their practices on data and AI are unsurprisingly limited. Multiplied by the gaps in their upskilling efforts, unauthorized AI use happens at least occasionally to 56% of organizations. Lack of effective organizational oversight or intentionality in how it is using AI compounds the deeper problem of agentic AI, where no human oversight may exist before an AI decision is made. This deepens the existing problematic status quo in which human end-users who often have little to no control over how their data is being used as a condition of using a service.

The same study confirms that 46% of organizations don't have a human-in-the-loop process for autonomous workflows, and 70% of leaders admit they're deploying agents faster than governance can keep up. Yet Canadian businesses who ban the use of AI may be adopting an unrealistic “out of sight, out of mind” solution, in which employees who wish to benefit from AI will quietly use it anyway. What is needed is serious governance guidelines that organizations can follow. If we want to uphold public trust and accountability of organizations, we must **establish a robust reporting regime**, and make their failure of compliance to reporting requirements consequential.

Disclosure on the use of agentic AI should be mandatory when its implementation directly influences or impacts the financial, legal or life costs or opportunities of any individual Canadian. When AI agents are deployed to interact with a human, the organization should be required to declare the use of AI agents and provide an alternative human contact point. When AI agents are positioned as a personal assistant or counsel for an individual, they must not be permitted to sell to or attempt to influence their users in any way. On top of this, **a right to appeal AI decisions must be guaranteed**. If organizations deploy AI agents to make important structural decisions, such as education, hiring, tenancy, or law enforcement, they must be required to perform independent systematic analysis for bias in their outputs, correct or discontinue their use as necessary.

We do not believe current efforts to improve AI transparency will hold as the final word on these complex and extraordinarily quickly changing systems; but reasonable efforts to update all our governance systems to account for them should begin now. The current transparency environment around agentic AI systems is already clearly insufficient, from technical solutions, to market practices and legal frameworks. We urge the government to **update consumer protection and contract law to clarify liability**, require disclosure of the use of AI, and provide a right of human appeal. This powerful technology also possesses the potential and capacity to undermine our democracy, therefore, we urge the government to ban the use of agentic AI or other AI systems to contact or otherwise target voters in Canadian electoral law with stiff penalties for violations.

Broader Considerations

Getting AI transparency right will likely take several bites at the regulatory apple. There is a great deal to study and iterate on here, which over time will form a more complete governance blanket. We'll close with a few of the most important starting points in our view.


First, AI transparency matters most when AI is involved in a consequential decision affecting a person's life, and when a human might reasonably mistake an AI system for a real human. These are the moments where the absence of disclosure causes the most harm, and where obligations should be concentrated first. **Obligations should scale with both risk and role**, rather than applying uniformly across all organizations. For example, small and medium-sized enterprises deploying low- to medium-risk AI should face lighter requirements than larger firms or high-risk deployments. Where risks are very high, this must be the predominant consideration.

Second, Canada should not embrace the Trump administration's foolhardy stance on AI. To maintain coherence with other jurisdictions, **Canada should pursue interoperable standards with the EU and OECD rather than the United States**, whose current AI governance approach is being shaped largely by industry preference, a so-called laissez-faire regulatory posture that is heavily driven by the Trump administration. Aligning with frameworks built around public accountability, rather than one built around minimizing obligations on industry, will better serve the goals this consultation has set out. And as noted earlier, Canada's deepening relationship with the EU offers a practical starting point for that alignment.

Finally, this consultation's focus is primarily on transparency, and we agree that reporting is the right starting point. Setting a high standard here should strengthen industry confidence rather than depress investment, since predictable, well-documented obligations reduce uncertainty for businesses more than they burden it. But AI technology will continue to evolve at an exponential rate, and any resulting regulation should not be static. Therefore, we urge the government to **build in a mandatory two-year review period into AI regulation** that includes public consultation to assess how the technology and its risks have changed since the framework's passage. The mandatory review and the ongoing administration of the framework should be in the hands of a **strongly independent regulator** that is insulated from government bodies that are simultaneously tasked with the promotion of AI investment, so that oversight is not compromised by a competing mandate to encourage adoption.

We thank you for the opportunity to share our thoughts on these matters, and are available for further discussion as useful for your next steps.

Sincerely,

A handwritten signature in black ink that reads "Matt".

Matt Hatfield
Executive Director, OpenMedia
matt@openmedia.org

A handwritten signature in black ink that reads "Jenna".

Jenna Fung
Policy Engagement Manager
jenna@openmedia.org