Advancing AI Transparency in Canada

Submission of the BC Freedom of Information and Privacy Association to Innovation, Science and Economic Development Canada

September 2026

Recognition

FIPA extends its sincere gratitude to the groups and individuals who generously provided valuable feedback on early versions and drafts, helping us refine this document. Many have also made their own submissions. Their combined input and expertise consistently provide important insights that strengthen our arguments for progressive law reform.

About FIPA

The BC Freedom of Information and Privacy Association (FIPA) is a non-partisan, non-profit society that was established in 1991 to promote and defend freedom of information and privacy rights in Canada. While we are based in BC, our membership extends across Canada, and we regularly partner with organizations throughout the country.

Our goal is to empower citizens by increasing their access to information and their control over their own personal information. We serve a wide variety of individuals and organizations through programs of public education, public assistance, research, and law reform. We are one of very few public interest groups in Canada devoted solely to the advancement of freedom of information and privacy rights.

Institutional Funders

The BC Freedom of Information and Privacy Association thanks the Law Foundation of BC, the Province of British Columbia (Gaming Policy and Enforcement Branch), and all our contributors including donors, funders, and volunteers for their ongoing support of our advocacy, programs, projects, and activities.





Framing

FIPA is not an AI organization, and this submission does not take a position on whether AI should be adopted. Our concern is narrower and follows from our mandate: we defend the public's right to know what public bodies are doing, and each person's right to control information about themselves. The adoption of AI affects both.

Personal information is a particularly valuable resource for the training and use of AI technologies, and current law and policy are inadequate to protect Canadians' personal data in an AI-powered data environment. These systems run on data, which gives every organization that adopts one a reason to collect more of it and to hold onto it longer.

When a public body or a regulated organization replaces a human process with an automated one, the transparency and privacy protections built around the human process do not carry over automatically. Records that used to exist stop existing. Notice that used to be given stops being given. Exemptions meant to be applied narrowly come to be applied to entire systems. None of this requires anyone to decide to reduce accountability. It happens as a side effect of changing how the work is done.

This is why we support placing transparency at the center of the AI governance discussion. We also want to be clear that transparency is foundational and that it is not sufficient on its own. Canada has been here before. Privacy law was built substantially around notice and consent, on the premise that if an organization tells you what it intends to do with your information and you agree, your interests are protected. What followed was a generation of technologies designed to extract consent rather than earn it, and the normalization of commercial surveillance at scale. Disclosure did not prevent that outcome, because someone who is told what is happening to them still has no ability to change it.

Transparency remains essential. It simply must be paired with substantive limits: on what can be collected, on which decisions can be automated, on who can obtain access to Canadians' information, and on what recourse a person has when an automated decision affects them. We encourage the government to treat the measures canvassed in the discussion paper as a foundation to build on rather than a complete response.

One dimension of this must be treated as core rather than as an afterthought. Transparency tells a person what is being done with their information; it cannot tell them who else is able to reach it. When the systems that process Canadians', information are owned, operated, or controlled by entities based in jurisdictions outside Canada, a foreign government may be able to compel access to that information regardless of where it is physically stored, and no Canadian disclosure requirement can reliably detect or prevent that access. Data and digital sovereignty is therefore not an adjacent trade or industrial policy question. It is a precondition of meaningful AI transparency, and we address it as such throughout this submission and directly in Recommendations 11 and 12.

Summary of Recommendations

1. A duty to inform individuals when their personal information is processed using AI
2. Deployers must disclose that a user is interacting with an AI system, in unambiguous terms
3. A general duty to label AI-generated content
4. Developers should be responsible for enabling detection of AI-generated content
5. A narrow approach to confidential business information
6. Model cards should be accompanied by a mandated plain-language summary
7. Mandatory incident reporting, with defined timelines
8. Defining a serious AI incident
9. Documentation and human-review standards for agents in high-risk contexts
10. Federal measures should be a floor, not a ceiling
11. Data and digital sovereignty must be built into AI transparency, not added afterward
12. Canada should renegotiate the trade commitments that constrain its ability to protect Canadians' data
13. Indigenous data sovereignty and the duty to consult
14. Employers should be required to publish AI policies and conduct PIAs before implementing AI systems
15. Extend AI transparency obligations to federal political parties

1. AI Interaction

Responding to consultation questions 5–8.

Recommendation 1 — A duty to inform individuals when their personal information is processed using AI

Public and private bodies should be required to inform individuals when their personal information is processed using an AI system.

This already exists provincially in a limited form. Quebec's *Act respecting Access to documents held by public bodies and the Protection of personal information* and Alberta's *Protection of Privacy Act* impose versions of this requirement on public bodies.

Bill C-36 moves partway there for the private sector: it would require an organization to explain an automated decision on request where that decision has a legal or similarly significant effect, and to give a general account of its use of automated decision systems.

Notification is not a courtesy. It is the precondition for every other right a person has over their own information. A person who does not know that an AI system processed their information cannot

request access to it, cannot ask for it to be corrected, and cannot complain about how it was used. Rights that depend on the individual to initiate them are only as good as the individual's knowledge that something has happened.

Automated processing raises the stakes in a way that ordinary record keeping does not. These systems do not only hold what a person provided; they generate new information about them by inference, and that inferred information can drive consequential decisions while never having been disclosed, verified, or corrected. Because AI systems depend on the volume and quality of data available to them, every organization that adopts one acquires a standing reason to collect more information and to retain it longer. A duty to inform is the minimum check on that dynamic.

The notification should be specific enough to be useful. At a minimum it should state:

- the categories of personal information being collected or used;
- the purpose or purposes of the AI processing;
- whether the processing occurs in-house or through a contractor or third-party platform;
- the jurisdiction or jurisdictions in which the information is stored, processed, or accessible, including whether the provider is subject to foreign laws that could compel its disclosure; and
- whether the output is a decision or a recommendation to a human decision-maker.

"This service uses AI" tells a person nothing they can act on.

We are asking for transparency as to both inputs and outputs. Input transparency covers what personal information enters the system and where it came from. Output transparency covers what the system produced, whether that is generated content or a decision affecting the person. Either one alone leaves a gap: knowing what went in tells a person nothing about what was concluded, and knowing the outcome tells them nothing about whether it was based on accurate information about them.

Recommendation 2 — Deployers must disclose that a user is interacting with an AI system, in unambiguous terms

Those who deploy AI systems should be responsible for informing users that they are interacting with AI.

The discussion paper notes the EU approach: providers must ensure that systems intended to interact directly with users are designed so that users are informed they are interacting with an AI system, unless it is obvious from the context. It also documents the problem clearly in its own case study: a Toronto man negotiated a car buy-back with a dealership chatbot named "Quinn" that never indicated it was not a person, and the dealership then attempted to revoke the offer on the basis that the chatbot had erred.

1. **Terminology matters and should be prescribed.** The paper acknowledges that some organizations disclose using softer language, such as "virtual assistant" or "built-in

intelligence", that leaves users no better informed. A disclosure requirement that permits this language does not accomplish anything. The term "artificial intelligence" should be required explicitly.

2. **Where the obligation sits.** The paper asks whether the onus should fall on developers or deployers. Our view is that it should fall on deployers, because deployers control the deployment context and the interface. Developers should bear a parallel obligation to build systems capable of making the disclosure.

2. AI-Generated Content

Responding to consultation questions 1–4.

Recommendation 3 — A general duty to label AI-generated content

Public and private bodies should be required to include disclaimers on AI-generated content.

Bill C-34, the *Safe Social Media Act*, would enact a *Digital Safety Act* requiring operators of regulated social media services to take adequate measures to label synthetic content: audio or visual representations of persons, objects, places, entities, or events that are likely to be mistaken for authentic recordings.[1] This is narrower than what we propose in two ways. It applies only to operators of regulated services, not to the developers or deployers of the systems that generate the content. And it reaches only content likely to be mistaken for something authentic, not AI-generated content generally.

We should be candid that the paper raises a real objection here: ubiquitous labelling may cause users to discount the signal, or to assume unlabelled content is authentic by default. Our response is that this is an argument for designing labels well, not for labelling narrowly. A label that appears only where content is deceptive requires someone to have already made the judgment call that the content is deceptive.

Recommendation 4 — Developers should be responsible for enabling detection of AI-generated content

Developers should bear responsibility for supporting detection of content their systems generate. Marking at the point of generation is the only stage at which the origin of content is known with certainty; downstream platforms are inferring.

The *Voluntary Code of Conduct on the Responsible Development and Management of Advanced Generative AI Systems* already recommends that developers establish a reliable and freely available detection method. Our position is that a voluntary recommendation is not sufficient where the whole detection ecosystem depends on it.

[1] Bill C-34, An Act to enact the Digital Safety Act and the Digital Safety Commission of Canada Act and to make consequential amendments to other Acts (Safe Social Media Act), 1st Sess, 45th Parl, 2026 (first reading 10 June 2026).

3. Information about AI Systems

Responding to consultation questions 9–12.

Recommendation 5 — A narrow approach to confidential business information

Question 10 asks directly how the government should balance increased transparency, including input transparency, against the protection of confidential business information. Our answer: confidential business information should be construed narrowly wherever a system is deployed for a public purpose or affects individual rights.

Government largely procures AI rather than building it. If vendors can assert commercial confidentiality over how a procured system works, procurement becomes a route to reduced transparency; not by anyone deciding to reduce it, but as a side effect of who owns what. FIPA has previously recommended increasing transparency in government procurement in the context of the *Access to Information Act*, and the same logic applies with more force here.

The concern is sharper where the vendor is controlled from outside Canada. A public body that procures an AI system from a provider subject to foreign law may be unable to say with confidence who else can access the information that system processes. Procurement disclosure should therefore include the vendor's ownership and control, the jurisdictions in which data is stored and processed, and any foreign legal obligations that could require the vendor to disclose Canadian data.

Confidentiality claims are not the only barrier. Article 19.16 of the Canada–United States–Mexico Agreement (CUSMA) prohibits Canada from requiring access to the source code of software, or to an algorithm expressed in that source code, as a condition of its sale or use, subject to a limited exception for a specific regulatory investigation or judicial proceeding.[2] A commitment of that kind sits uneasily with any transparency regime that depends on examining how an AI system works before harm occurs. We return to this in Recommendation 12.

Recommendation 6 — Model cards should be accompanied by a mandated plain-language summary

The paper describes model and system cards as the main market response, and notes that the growing volume and technical complexity of model cards can make it difficult for the public to assess them. That matches a concern FIPA has heard consistently: model cards are too complicated for a general audience.

Rather than replacing model cards, we recommend requiring a short, standardized, plain-language summary alongside them, covering intended purpose, known limitations, whether personal

[2]Canada–United States–Mexico Agreement, Chapter 19 (Digital Trade), arts 19.2, 19.11, 19.12 and 19.16, online: Global Affairs Canada https://www.international.gc.ca/trade-commerce/trade-agreements-accords-commerciaux/agr-acc/cusma-aceum/text-texte/19.aspx?lang=eng.

information was used in training, what human oversight applies in deployment, and where the system and the data it processes are hosted and under which jurisdiction's law.

4. AI Incidents

Responding to consultation questions 13–16.

Recommendation 7 — Mandatory incident reporting, with defined timelines

Canada should institute incident reporting obligations for providers of AI systems, drawing on two existing models:

- the EU's *General Data Protection Regulation* (GDPR), which requires reporting of breaches involving personal information within 72 hours; and
- Article 73 of the EU AI Act, which requires providers of high-risk systems to report serious incidents to the market surveillance authority in the country where the incident occurred, as soon as a causal link is established or reasonably suspected: within 15 days generally, two days for widespread or serious incidents, and ten days where a death has occurred. Initial incomplete reports are permitted if followed by a full report, and providers must investigate and cooperate with authorities.

Our recommendation is a combination: the EU's tiered structure for AI incidents generally, with a 72-hour period where personal information is involved, and a shorter period where sensitive personal information is involved.

Recommendation 8 — Defining a serious AI incident

The consultation asks how a serious AI incident should be defined. The EU definition covers death or serious harm to health, serious and irreversible disruption to critical infrastructure, infringement of fundamental rights obligations, and serious harm to property or the environment.

From FIPA's perspective, the fundamental-rights limb is the one that matters most and the one most likely to be dropped in a Canadian adaptation. A Canadian definition should expressly include:

- unauthorized disclosure, exposure, or inference of personal information by or through an AI system;
- disclosure of personal information processed by an AI system to a foreign government or authority, including disclosure compelled under foreign law;
- automated decisions affecting a person's rights, benefits, or employment that turn out to have been based on erroneous processing; and
- failures in systems used by public bodies to deliver services or process requests, including access and privacy requests.

5. AI Agents

Responding to consultation questions 17–20.

Recommendation 9 — Documentation and human-review standards for agents in high-risk contexts

AI agents deployed in high-risk contexts should operate under specific standards for documentation and human review. The EU AI Act does not create a separate category for agents. An agentic system is regulated according to how it is used: where it is deployed in one of the high-risk contexts listed in Annex III, the Act's high-risk obligations apply to it, including automatic record-keeping (Article 12) and human oversight (Article 14). The Act's provisions on general-purpose AI models govern the underlying model, not the agentic system built on it. Annex III provides a workable starting list of high-risk contexts for Canada to adapt.

The problem is one FIPA raised in our submission on reform of the *Privacy Act*: agentic systems make the precise path of a decision very difficult to reconstruct. That is a direct problem for rights that already exist, including the right to access how a decision was made and the right to correction. A right of access to reasoning that no one recorded is not a right.

The consultation's own material supports this. It notes that as agents interact with other agents, it becomes increasingly difficult to identify which agent did what, when, and why. It describes an emerging market in observability tooling, including agent platforms that log actions, track tool usage, and record decision pathways, and standardization work such as Open Telemetry's semantic conventions for AI systems. It also observes that some businesses have taken the position that users bear responsibility when agents do not behave as directed.

That combination is the argument. The technical capacity to log agent activity already exists and is commercially available. What is missing is any obligation to use it, and any entitlement for an affected person to see the result.

Where AI agents operate in high-risk contexts, we recommend that:

- action-level logs be required, recording what was invoked, when, with what parameters, and under whose authority.
- logs be retained long enough to survive a complaint or review timeline, because a log deleted before a complaint is resolved is not an accountability measure.
- meaningful human review be required at defined decision points, rather than merely available in principle; and
- where an agent acts affecting a person, that person be able to obtain the record of what it did.

6. Crosscutting

Responding to consultation questions 21–24.

Recommendation 10 — Federal measures should be a floor, not a ceiling

Question 23 asks how Canada should ensure coherence between federal measures and provincial, territorial, and international frameworks.

Coherence is a reasonable goal, but harmonization in practice often means levelling to the least demanding common standard, because that is the one all parties can accept. Several provinces are already ahead of the federal government: Quebec and Alberta on notice for automated processing of personal information, and Ontario on AI disclosure in hiring and on public sector AI accountability frameworks under the *Enhancing Digital Security and Trust Act, 2024*.

Federal measures should be expressly framed as minimum standards, leaving provinces free to legislate more protective requirements, and any harmonization language should be reviewed specifically for whether it would displace existing provincial protections.

Recommendation 11 — Data and digital sovereignty must be built into AI transparency, not added afterward

Question 21 asks whether there are other areas that should be considered for action to enhance AI transparency. Data and digital sovereignty are the most significant gap in the discussion paper.

The term now carries several meanings. Teresa Scassa, Canada Research Chair in Information Law and Policy, observes that sovereignty is invoked to describe state control over data and technology, Indigenous peoples' control over data as an element of self-determination, and individuals' control over information about themselves.[3] Each is engaged by AI. For the purposes of this submission, the point is a practical one: a transparency obligation is only as meaningful as the ability of Canadian law to reach the data and systems it describes.

That ability is not secured by physical location. A senior Microsoft France official testified before the French Senate that, under the U.S. *CLOUD Act*, the company could be compelled to hand over data from any country regardless of where it is stored; France has since announced it will move government videoconferencing off U.S.-based platforms, and other European states are pursuing alternatives.[4]

Canadian data and digital sovereignty therefore require more than enabling foreign companies, subject to foreign laws and to the influence of foreign governments, to build and manage digital infrastructure in Canada. A data centre on Canadian soil, operated by a company that can be compelled to hand its contents to another government, does not protect Canadians. Sovereignty requires independent domestic ownership, management, and oversight of the infrastructure that holds Canadians' most sensitive information, and that requirement should be treated as non-negotiable.

---

[3]Teresa Scassa, "Sovereignty and the Governance of Artificial Intelligence" (2024) 71 UCLA L Rev Disc 214, online: https://www.uclalawreview.org/wp-content/uploads/securepdfs/2024/05/06-Scassa-No-Bleed.pdf.

[4]Julie Sobowale, "Canada's uncertain digital future", CBA National (20 April 2026), online: https://nationalmagazine.ca/en-ca/articles/in-depth/2026/canada-s-uncertain-digital-future.

The stakes have grown. Canadian legal commentators note that Canada long assumed a stable relationship with the United States and can no longer do so, and they identify foreign interference in Canada's democratic processes as a real and significant threat.[5] Canadians' personal information must be protected from misuse by foreign and domestic actors alike, and that protection cannot depend on the continuing goodwill of any foreign government. Where the companies that hold Canadians' data can be directed or compelled by a foreign state, the distinction between the intentions of allied and adversarial governments offers little practical assurance.

This is a transparency issue as much as a security issue. A person told that "this service uses AI" is not told that their information is held by a company answerable to another country's surveillance laws. We recommend that:

- organizations subject to a federal AI transparency obligation be required to disclose, as part of the information provided under Recommendations 1 and 6, the jurisdictions in which personal information is stored and processed, the ownership and control of the entities that store and process it, and whether those entities are subject to foreign laws that could compel its disclosure;
- AI systems used by federal public bodies to process personal information, or to make or support decisions about individuals, be hosted and operated on infrastructure that is Canadian-owned, Canadian-managed, and subject to Canadian oversight, with any exception approved at a senior level and published.
- organizations be required to report, at least in aggregate, the foreign government demands they receive for Canadians' personal information, and to notify affected individuals and the relevant Canadian regulator wherever doing so is not prohibited; and
- the government treat sovereignty as a core design requirement of any national data strategy and of public digital infrastructure, rather than as a procurement preference that can be traded against cost.

Indigenous data sovereignty is distinct from national digital sovereignty and is addressed in Recommendation 13. The two should not be conflated, and national sovereignty measures must not be used to override Indigenous governance of Indigenous data.

Recommendation 12 — Canada should renegotiate the trade commitments that constrain its ability to protect Canadians' data

Question 23 asks how Canada should ensure coherence between federal measures and international frameworks. Coherence must run in both directions: Canada's trade commitments should not prevent it from adopting the transparency and sovereignty measures recommended in this submission. Digital sovereignty has become a live issue in Canada's trade discussions,[6] and it should be addressed as such.

[5]Sobowale, supra note 4.

[6]CBC Radio, "Digital sovereignty has come up in trade talks, but what does it mean?", Just Asking, online: https://www.cbc.ca/radio/justasking/canada-digital-sovereignty-just-asking-9.7337778.

Chapter 19 (Digital Trade) of CUSMA prohibits Canada from restricting cross-border transfers of information, including personal information, where the transfer is for the conduct of a covered person's business (Article 19.11); from requiring a covered person to use or locate computing facilities in Canada as a condition of doing business here (Article 19.12); and, as noted in Recommendation 5, from requiring access to source code or algorithms as a condition of sale or use (Article 19.16).[7] Article 19.11 permits measures necessary to achieve a legitimate public policy objective, but only on conditions that a trade panel would ultimately assess. Critics have long observed that these disciplines, which build on the e-commerce chapter of the *Comprehensive and Progressive Agreement for Trans-Pacific Partnership* (CPTPP), give primacy to the unimpeded flow of data across borders while narrowing governments' room to regulate the digital economy in the public interest.[8]

This is not a new concern for FIPA. A decade ago, records obtained by FIPA revealed a steady stream of meetings, memos, and negotiating pressure on Canada to weaken data residency rules, and the U.S. Trade Representative listed the protections in British Columbia's *Freedom of Information and Protection of Privacy Act* as a foreign trade barrier.[9] British Columbia relaxed those protections in 2021. Domestic data protection has been treated as a bargaining chip in trade negotiations, and Canada has conceded more than it should have.

The timing matters. At the first joint review of CUSMA on July 1, 2026, the United States declined to confirm a sixteen-year extension, stating that it would not renew the Agreement in its current form. The Agreement remains in force, but it is now subject to annual joint reviews until the parties agree to extend it or it expires in 2036.[10] Canada has an opportunity it has not had since 2018 to reopen these provisions. We recommend that the government:

- use the CUSMA annual review process to seek the renegotiation or removal of Articles 19.11, 19.12, and 19.16, or at minimum express carve-outs for measures that protect personal information, require domestic hosting of sensitive data, or give Canadian regulators access to AI systems for transparency and oversight purposes;
- state clearly that Canada regards the protection of Canadians' personal information from foreign access, and the integrity of its public digital infrastructure, as essential security interests within the meaning of Article 32.2, which preserves each party's ability to apply measures it considers necessary to protect its essential security interests;[11]

[7]CUSMA, Chapter 19, supra note 2.

[8]Stuart Trew, "CUSMA 2026: Reversing Big Tech's power grab", Canadian Centre for Policy Alternatives (22 August 2024), online: https://www.policyalternatives.ca/news-research/cusma-2026-reversing-big-techs-power-grab/.

[9]Scott Sinclair, "TPP Deal Puts BC's Privacy Laws in the Crosshairs", Canadian Centre for Policy Alternatives (16 July 2015), online: https://www.policyalternatives.ca/news-research/tpp-deal-puts-bcs-privacy-laws-in-the-crosshairs/; see also FIPA, "US Trade Representative calls BC privacy law a trade barrier", online: https://fipa.bc.ca/us-trade-representative-calls-bc-privacy-law-a-trade-barrier/.

[10]White & Case LLP, "USMCA 2026 Joint Review: United States declines to extend Agreement, triggering annual reviews" (2 July 2026), online: https://www.whitecase.com/insight-alert/usmca-2026-joint-review-united-states-declines-extend-agreement-triggering-annual; Canada–United States–Mexico Agreement, art 34.7.

[11]Canada–United States–Mexico Agreement, art 32.2 (Essential Security).

- preserve in any renegotiation the exclusion of information held or processed by or on behalf of government from the scope of Chapter 19 (Article 19.2(3)); and
- apply the same approach to the CPTPP and to future trade agreements, so that AI transparency and data protection measures are not negotiated away as trade concessions.

We recognize that reliance on Article 32.2 in this context is untested and that invoking it carries diplomatic cost. That is precisely why Canada's position should be stated now, before a dispute arises, rather than improvised in response to one.

Recommendation 13 — Indigenous data sovereignty and the duty to consult

Any federal AI transparency framework should reinforce existing constitutional obligations, and Canada's commitments under the *United Nations Declaration on the Rights of Indigenous Peoples*, to consult with First Nations, Inuit, and Métis peoples.

FIPA is not an Indigenous organization and does not speak for Indigenous nations. We note, however, that AI systems trained on or applied to data about Indigenous communities raise questions a general transparency framework will not answer on its own. The First Nations Information Governance Centre is the custodian of the First Nations principles of OCAP® (ownership, control, access, and possession), and Inuit and Métis peoples have their own distinct approaches to data governance.[12] As Scassa notes, Indigenous data sovereignty concerns far more than where data is stored; it concerns the right to govern the collection, ownership, and application of data about Indigenous communities, peoples, lands, and resources.[13]

Indigenous governments and organizations should be consulted on how AI transparency obligations apply to data about their communities, and national data sovereignty measures should be designed to support, rather than displace, Indigenous governance of Indigenous data. At the same time, as FIPA noted in our submission on the *Access to Information Act*, measures framed as protective must not become a basis for withholding information that should be released.

Recommendation 14 — Employers should be required to publish AI policies and conduct PIAs before implementing AI systems

Employers should be required to publish AI policies and to conduct privacy impact assessments (PIAs) before implementing AI systems that process employees' personal information. For federally regulated private-sector employers, this could be achieved through PIPEDA or its successor legislation; for federal public sector employees, through the *Privacy Act*.

Policies published for employees and the public should explain how AI is being used and for what purpose, whether personal information will be included, whether the system will handle sensitive information, how AI incidents will be mitigated, and, where applicable, how the results will be used

[12] Sobowale, supra note 4.

[13] Scassa, supra note 3.

in decision-making. These elements are consistent with recommendations made by Québec's privacy regulator, the Commission d'accès à l'information (CAI).

PIAs are a practical tool for estimating risk before harm occurs. They should expressly address algorithmic bias and the potential for discriminatory outcomes.

Recommendation 15 — Extend AI transparency obligations to federal political parties

Federal political parties (FPPs) routinely manage large volumes of personal information, including sensitive information about political opinions, and may use AI systems to do so. AI transparency obligations should apply to them as they would to any other organization. The content FPPs produce is also uniquely influential, and AI-generated political content should be subject to the labelling duty in Recommendation 3.

This is consistent with FIPA's longstanding recommendation that FPPs be brought under federal privacy law. Bill C-25, the *Strong and Free Elections Act*, would prohibit the creation of deepfakes of candidates, other than for satire or parody.[14] That is a welcome but narrow step: it addresses the most deceptive form of synthetic political content without establishing any general transparency obligation for parties' own use of AI.

[14]Bill C-25, Strong and Free Elections Act, 1st Sess, 45th Parl, 2026; see Sobowale, supra note 4.

Appendix 1: ISED Survey Questions

1. Would it help you trust what you see online if you could tell whether something was created by AI or by a person? When and for what kinds of content (e.g., image, video, audio) is it most important to know if something was created or modified by AI?
2. What would best help Canadians determine when content is AI-generated (e.g., invisible or visible watermarks, disclaimers, provenance metadata)?
3. Who in the AI value chain (developers, deployers, or others) should be responsible for providing transparency around AI-generated content? Why?
4. Do existing market practices, technical tools, and legal frameworks make it easy enough to know when content is AI-generated? If not, where do gaps remain and what actions (e.g., regulatory measures, guidance and codes of conduct, standards and technical solutions, research and development, literacy initiatives, procurement requirements) do you think the Government should take?
5. When and why is it most important to know that you are interacting with an AI system? Are there contexts where you don't need to know that you're interacting with AI?
6. What factors make a disclosure of AI use meaningful and effective for the user?
7. Who in the AI value chain (developers, deployers, or others) should be responsible for providing transparency around AI interactions? Why?
8. Are existing market practices and legal frameworks sufficient to support transparency around AI interactions? If not, where do gaps remain and what actions (e.g., regulatory measures, guidance and codes of conduct, standards and technical solutions, research and development, literacy initiatives, procurement requirements) do you think the Government should take?
9. What kinds of information about AI would you like to know in order to help you make informed choices about when and how to use AI? When and for which types of AI products is this information most important to you?
10. What kind of information about AI do businesses need to adopt AI safely and responsibly? Do businesses currently have access to the information that they need? How can the Government best balance the need to increase transparency (including input transparency) against the need to protect confidential business information?
11. Who in the AI value chain (developers, deployers, or others) is best placed to provide what information about AI systems, how they are created, and their capabilities and limitations?
12. Are existing disclosure practices and legal frameworks sufficient to provide the information that users and other stakeholders need? If not, where do gaps remain and what actions (e.g., regulatory measures, guidance and codes of conduct, standards and technical solutions, research and development, literacy initiatives, procurement requirements) do you think the Government should take?

13. Do the public and the Government have enough information about AI risks, incidents, and potential future harms? What types of information related to AI use and risk trends would be most useful so that Canadians and the Government can proactively address AI-related risks and harms?
14. How should a serious AI incident be defined, and what information would be most useful to better understand the causes and consequences, as well as the mitigation, of these incidents? How could reporting requirements be structured to encourage proactive disclosure of AI incidents and also address concerns around confidentiality and liability?
15. Which individuals or actors (e.g., developers, deployers) are best placed to provide the most relevant information? Who (e.g., regulatory bodies, nonprofits, others) should collect this information and what should be done with it?
16. Are existing incident reporting frameworks in Canada (e.g., for consumer products, medical devices, transportation) sufficient to address serious incidents involving AI? If not, where do gaps remain and what actions (e.g., regulatory measures, guidance and codes of conduct, standards and technical solutions, research and development, literacy initiatives, procurement requirements) do you think the Government should take? Are there emerging international best practices that Canada could adapt or interoperate with?
17. Are you concerned about the use of AI agents? In what contexts is it most important to know when an AI agent is in use? Why?
18. Who should be responsible for disclosing the use of AI agents? What kinds of information should they be recording regarding the activities of AI agents, and what kinds of disclosures should they be making, and to whom?
19. What kinds of information do businesses need in order to confidently and safely use AI agents? Do businesses currently have access to the information that they need about AI agents?
20. Are existing technical solutions, market practices, and legal frameworks sufficient to support transparency around agentic AI systems? If not, where do gaps remain and what actions (e.g., regulatory measures, guidance and codes of conduct, standards and technical solutions, research and development, literacy initiatives, procurement requirements) do you think the Government should take?
21. Given the areas discussed here, when is AI transparency most important, and are there any areas where action would be premature? Are there other areas that should be considered for action to enhance AI transparency?
22. How should any measures account for the circumstances of small and medium-sized enterprises?
23. How should Canada ensure coherence between any federal measures and provincial, territorial, and international frameworks? What approaches in other jurisdictions could best inform a Canadian consideration of this issue?

24. How can Canada design measures to best account for the ongoing evolution in AI technologies and their use?