
CPA Submission to Innovation, Science and Economic Development Canada re: “Have your say on advancing AI transparency in Canada”

On July 23rd, 2026 Innovation, Science and Economic Development Canada opened a consultation for the public to submit their feedback on Canada’s National Artificial Intelligence Strategy: AI for All. As part of the CPA’s commitment to being the national voice and thought leader for psychology in Canada, we made a submission to this public consultation using the survey option that was provided as part of the consultation.

The consultation was primarily focused on the broad theme of improving AI transparency in Canada and so our responses were closely aligned with this theme. Among the themes that we emphasized in our submission were:

- 1) A need to strengthen the ability to distinguish AI content from human-generated content, with clearly visible watermarks that are difficult for end users to remove and adhere to strong provenance (i.e., the record of origin of content) requirements.
- 2) The importance of AI interaction disclosure and protections, with a particular emphasis on the importance of this disclosure in sensitive health information contexts.
- 3) A need for clear data capture and retention policies that are written in plain language, are easily accessible, and make explicit how and why data are being used. We emphasized in particular the importance of not allowing data policies to be obscured by physical server locations when they interact with foreign data laws.
- 4) The importance of clear reporting structures for new AI information, including publication, reporting, and action for serious AI incidents.
- 5) The additional challenges and need for enhanced protections when deploying and interacting with autonomous AI agents, with a particular emphasis on the need for agent oversight requirements and clear abilities to monitor and intervene when issues arise.

These were just some of the themes that were covered in our submission. If you are interested in reading our submission in greater detail you can review our entire submission appended at the end of this document.

Appendix A: Public Consultation Survey Questions & Answers

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?

Yes. Clear identification of AI-generated or AI-modified content would help Canadians assess its source, purpose, reliability, and potential risks. This matters for all content types including text, images, video, and audio because each can mislead or materially shape how people understand an issue. Disclosure is particularly important for content involving health or mental health, public safety, news, elections, financial or legal decisions, education, employment, public benefits, or any other context in which inaccurate or deceptive information could cause significant harm. The disclosure should distinguish between content created entirely by AI and content that was only edited or supplemented by AI, and should indicate which portions were affected where feasible.

What would best help Canadians determine when content is AI-generated (e.g., invisible or visible watermarks, disclaimers, provenance metadata)?

The most useful approach would combine a permanent, readily visible label or watermark with tamper-resistant metadata. The visible disclosure should tell an ordinary user that AI was involved, whether the content was generated or modified, and, where feasible, which portions were affected. The metadata should provide more detailed information about when and how AI was used, the tools involved, and the content's subsequent modification history. These measures should follow common technical standards, remain attached when content is copied or shared, and be difficult for end users to remove. AI and standards experts should advise on technical feasibility and implementation.

Who in the AI value chain (developers, deployers, or others) should be responsible for providing transparency around AI-generated content? Why?

Responsibility should be shared across the AI value chain, but each actor should be accountable for the risks it is best positioned to control. Developers should build persistent, tamper-resistant labels, watermarks mechanisms into their systems and provide accurate technical information about how they work. Deployers should ensure that disclosures remain visible, accurate, and appropriate to the context in which the content is used. Platforms that distribute content should preserve watermark and other relevant information and apply consistent presentation rules. End users should not be able to remove required disclosures. Government should establish and enforce minimum standards, clarify responsibility when several actors are involved, and provide effective remedies when transparency requirements are ignored.

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?

No. Existing market practices, technical tools, and legal frameworks do not make it easy enough to identify AI-generated or AI-modified content. Detection is incomplete, labels are inconsistent, and disclosures often fail to explain the extent or nature of AI involvement. For example, a three-minute

video containing one two-second AI-generated segment may receive the same label as a video created entirely by AI. Users need disclosures that distinguish generated from modified content, identify the affected portions where feasible, and remain attached when content is shared. The Government should establish clear regulatory requirements, common technical and provenance standards, codes of conduct, public-sector procurement conditions, and public literacy initiatives. Because labels alone may not prevent misinformation or disinformation from spreading, these measures should be paired with research and evidence-based interventions addressing how misleading content is created, amplified, interpreted, and shared. AI, communications, psychology, and other relevant experts should help design and evaluate those interventions.

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?

It is most important to know that AI is being used when the system collects sensitive information, communicates directly with a person, makes or materially influences a decision, provides advice, or affects access to a service, benefit, opportunity, or right. This includes health and mental health care, employment, education, financial and legal services, public benefits, government services, and other moderate- or high-risk settings. In these contexts, people should be told what the AI does, what data it uses, the extent of human oversight, who is accountable, and how to seek review. They should also have a meaningful opportunity to opt out and use a non-AI alternative without reduced access, lower-quality service, or degraded care. Disclosure may be less important in genuinely low-risk uses (e.g., when AI merely assists a person in drafting a routine email that the person reviews and sends). The assessment should consider the sensitivity of the information, the AI system's level of autonomy, the degree of human oversight, and the potential consequences for affected people.

What factors make a disclosure of AI use meaningful and effective for the user?

A meaningful disclosure must be timely, prominent, understandable, accessible, and specific enough to support an informed choice. It should explain that AI is being used; it should indicate whether the AI generated, modified, recommended, or decided something; what data it collects or accesses; whether those data are retained or used for training; what decisions it can make or influence; the extent of human oversight; and who is accountable. It should also explain important limitations, risks, and available avenues for review, correction, complaint, or redress. In professional and other moderate- or high-risk settings, disclosure must be paired with a genuine opportunity to opt out and access a non-AI alternative without reduced service or care. Requiring a person to refuse the entire service is not meaningful choice or informed consent.

Who in the AI value chain (developers, deployers, or others) should be responsible for providing transparency around AI-generated content? Why?

Responsibility should be shared, with each actor accountable for the information and safeguards it controls. Developers should document system capabilities, limitations, data practices, risks, and implement technical safeguards, and should build persistent disclosure mechanisms into their systems. Deployers should explain how AI is used in the specific interaction, maintain appropriate human oversight, provide a clear accountable contact, and preserve records needed for audit or review. Platforms and intermediaries should preserve required labels and provenance information. Governments

should establish and enforce baseline requirements, while sectoral regulators and professional bodies should set additional standards for disclosure, consent, privacy, record-keeping, incident reporting, meaningful opt-outs, and non-AI alternatives in regulated or high-risk settings.

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?

No. Existing market practices and legal frameworks do not provide consistent or meaningful transparency around AI interactions. Requirements vary across sectors, disclosures are often vague or buried in terms of service, and users may not be told what data are used, whether those data are retained or used for training, what decisions the system can make or influence, how much human oversight exists, or how to seek review or redress. The Government should establish plain-language baseline disclosure requirements and stricter sector-specific rules for moderate- and high-risk uses. These should be reinforced through technical standards, codes of conduct, independent testing and verification, public literacy initiatives, research, and procurement requirements. Regulated professions should also adopt clear standards governing disclosure, informed consent, meaningful opt-outs, privacy, documentation, human oversight, incident reporting, accountability, and access to non-AI alternatives.

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?

To make an informed choice, users need clear, reliable, and independently verifiable information about a system's purpose, capabilities, limitations, accuracy, known risks, level of autonomy, human oversight, accountability, and data practices. They should be able to determine what data are collected, where they are stored and processed, whether they are retained or used to train or improve models, which third parties or foreign jurisdictions may access them, how long they are kept, and how they can be corrected or deleted. Claims that data are stored in Canada should also explain any legal or contractual arrangements that may permit access under foreign laws. Users should also know how to opt out, obtain human review, report a problem, and seek review. This information is especially important for products used with health, mental health, employment, financial, legal, educational, government, or other sensitive information, and for systems that make or materially influence decisions.

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?

Businesses need accurate, standardized information about an AI system's intended uses, capabilities, limitations, accuracy, known risks, security controls, data flows, model and vendor dependencies, human-oversight requirements, auditability, incident history, and applicable legal and professional obligations. They should be able to determine what information enters the system; where it is stored and processed; whether it is retained or used for training; which third parties or jurisdictions can access it; how outputs can be reviewed or challenged; and how the system can be stopped, corrected, or replaced.

Current information is often incomplete, inconsistent, or framed as marketing. The Government should require standardized disclosure of material risks and the full electronic information chain, supported by independent verification for higher-risk systems. Genuine trade secrets may be appropriate to receive narrowly tailored protection, but confidentiality should not shield information needed to assess safety, privacy, accountability, compliance, or harm.

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?

No. Existing disclosures are often incomplete, inconsistent, difficult to compare, or expressed in broad marketing language. Statements that a service is “made in Canada” or uses Canadian servers may not explain the full digital chain, including where data travel, who can access them, whether they are used for training, which laws apply, and how long they are retained. The Government should require standardized, plain-language disclosures covering purpose, capabilities, limitations, accuracy, risks, human oversight, accountability, data collection and access, retention, training use, third-party involvement, and available review or redress, as well as explicit requirements for this implementation in terms of use/service. Higher-risk systems should be subject to independent testing or verification, incident reporting, and sector-specific requirements. Codes of conduct, technical standards, research, public literacy initiatives, and procurement rules should reinforce these obligations, with public-sector procurement conditional on adequate transparency, privacy, security, and auditability.

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?

No. The public and the Government need a consistent, independent, and accessible source of evidence-based information about AI risks, incidents, and emerging harms. Useful reporting should cover privacy and security breaches, inaccurate or discriminatory outputs, harmful advice, misinformation and disinformation, failures of human oversight, misuse, near misses, and impacts in high-risk sectors. It should identify affected populations, recurring system or sector patterns, the severity and frequency of harms, corrective actions, and whether mitigations were effective. Canadians should also be able to learn what information is stored about them, how inaccuracies can arise, how systems are changing, and how those changes may affect their rights and daily lives. Public reporting should be anonymized where necessary, but sufficiently detailed to support research, regulation, procurement, public literacy, and prevention. Reliance primarily on information supplied by AI companies or repeated through news coverage is insufficient because it may be incomplete, inconsistent, or biased. Finally, there is a lack of public knowledge about the true environmental impacts of AI, which should be made more transparent for the public.

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?

A serious AI incident should include any event in which an AI system's output, action, omission, security failure, unauthorized use, or foreseeable misuse causes—or creates a credible risk of—death, physical or psychological injury, discrimination, loss of rights or essential services, significant financial or reputational harm, exposure of sensitive information, or substantial harm to the public. Serious near misses and recurring lower-level incidents should also be reportable when they reveal a systemic risk. Reports should identify the system and version involved; the developer and deployer; the context, affected population, and relevant inputs and outputs; the data accessed; the role of human oversight; the actual or potential harm; contributing factors; prior similar incidents; corrective action; and the effectiveness of mitigation. Developers and deployers should report serious incidents promptly, with protections for good-faith and whistleblower reporting, proportionate confidentiality safeguards, and liability rules that encourage early disclosure without removing accountability. Members of the public and professionals should also have a simple reporting route and receive confirmation that their report will be reviewed.

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?

Developers are best placed to provide technical information about system design, training, testing, safeguards, known limitations, model changes, and model-level incidents. Deployers are best placed to report the context of use, data accessed, human oversight, affected individuals, operational consequences, and corrective action. Platforms, vendors, and professional users may also hold relevant distribution, integration, or practice information. Each actor should be legally responsible for maintaining records and supplying the information within its control. An independent government body with regulatory and enforcement authority should receive and triage reports, coordinate with privacy, human rights, consumer, health, transportation, financial, professional, and other sectoral regulators, and refer matters where appropriate. It should publish anonymized trend information, identify recurring risks, support research and public literacy, issue guidance or corrective orders, inform procurement decisions, and use the information for compliance, prevention, and redress.

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?

Existing Canadian sectoral reporting frameworks may provide useful foundations, but they are not sufficiently clear, consistent, or integrated for serious AI incidents. From a user's perspective, it may be unclear which body is responsible, what threshold applies, what information must be provided, whether the report will be reviewed, or how incidents spanning several sectors will be handled. The Government should create a single, low-friction reporting entry point that routes matters to the appropriate regulator; adopt common definitions, severity levels, timelines, and minimum reporting fields; clarify the responsibilities of developers, deployers, platforms, and professional users; and provide feedback to reporters. Existing regulators should retain sector-specific expertise, supported by interoperable

technical standards, shared data, professional guidance, public literacy, independent oversight, and transparent accountability. Canada should also assess and align with emerging international standards and reporting models where interoperability would improve cross-border detection and response, while adapting them to Canadian privacy, human rights, health, consumer-protection, and professional-regulation requirements.

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?

Yes. AI agents create many of the same risks as other AI systems, but those risks are amplified when an agent can act autonomously, communicate externally, access sensitive information, initiate transactions, alter records, or make decisions without timely human review. Disclosure is especially important when an agent handles personal, health, financial, legal, employment, educational, or government information; makes or materially influences decisions; purchases goods or services; or communicates on behalf of a person or organization. Users should always know that they are interacting with an AI agent, what tasks and decisions it is authorized to undertake, what data and systems it can access, what safeguards and approval thresholds apply, when a human is involved, how the agent can be stopped or overridden, and who is accountable if a concern or incident arises. In moderate- and high-risk settings, users should also have access to human review and a meaningful non-AI alternative.

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?

The deployer should be primarily responsible for disclosing that an AI agent is in use because it determines the agent's role, permissions, and operating context. Developers and vendors should provide the technical information needed to make that disclosure accurate and should remain accountable for defects, undocumented limitations, and inadequate safeguards within their control. Users should be told that they are interacting with an agent; what tasks it can perform; what decisions it can make or influence; what data and systems it can access; what transactions or communications it may initiate; when human approval is required; how it can be stopped or overridden; and how to contact an accountable person. Deployers should maintain auditable records of the agent's instructions, permissions, material actions, decisions, external communications, data access, human interventions, errors, overrides, and incidents. Affected users should receive information relevant to their interaction, while regulators, auditors, and professional oversight bodies should have access to the more detailed records needed for oversight, subject to appropriate privacy, security, and confidentiality protections.

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?

Businesses need clear, independently verifiable information about an AI agent's capabilities, limitations, accuracy, autonomy, permissions, data access, security controls, model and vendor dependencies, interoperability, human-oversight requirements, auditability, incident history, and legal and professional obligations. They should know which tasks can safely be delegated, which decisions or transactions require human approval, how the agent can be monitored, stopped, overridden, corrected, or rolled back, and who remains accountable. They also need a complete view of where data travel, whether they

are retained or used for training, and which third parties or jurisdictions can access them. Current access to this information appears inconsistent and often depends on vendor claims. Standardized documentation, independent testing for higher-risk systems, audit logs, contractual clarity, and sector-specific guidance are needed. In health care and other regulated settings, agents should also comply with the profession's ethical and practice standards and provide meaningful human oversight and a non-AI alternative.

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?

No. Existing technical solutions, market practices, and legal frameworks do not provide consistent transparency or accountability for agentic AI systems. Key gaps include unclear disclosure; limited auditability; insufficient information about permissions, data access, autonomous actions, model and vendor dependencies, and system changes; uncertain responsibility across developers, vendors, deployers, platforms, and professional users; and inconsistent human-oversight, opt-out, and incident-reporting requirements. The Government should establish enforceable baseline rules for plain-language disclosure, documented permissions, data governance, security, testing, record-keeping, human approval and override, incident and near-miss reporting, independent audit, accountability, and redress. Higher-risk uses should face stricter sector-specific requirements, including meaningful human review and access to a non-AI alternative. Codes of conduct, interoperable technical and provenance standards, professional guidance, public literacy initiatives, research, independent oversight, and procurement requirements should reinforce these rules. Procurement should require evidence that an agent is transparent, secure, auditable, controllable, and compliant before it is used in public services.