



Enhancing Trust in Artificial Intelligence Through Increased Transparency

The Information & Communications
Technology Council's submission to the Government
of Canada's public consultation on AI transparency

September 23, 2026

About ICTC

The Information and Communications Technology Council (ICTC) is a national, neutral, not-for-profit centre of expertise whose mission is to strengthen Canada's digital advantage in the global economy. For more than 30 years, we have provided forward-looking research, practical policy advice, and innovative, industry-informed capacity development solutions for individuals, businesses, and the public sector. Comprised of a national team of experts from coast to coast, our goal is to ensure that technology is used to drive economic growth and innovation, and that Canada's workforce remains globally competitive.

Executive Summary

ICTC welcomes the Government of Canada's consultation on measures to improve, strengthen, and increase AI transparency. Clear and proportionate transparency requirements can strengthen public confidence, improve business adoption, support regulatory oversight, and give Canadians meaningful ways to understand and address the consequential uses of AI.

To support the Government of Canada in its efforts to advance AI transparency and create the regulatory, policy, and governance frameworks required to strengthen trust in AI, ICTC has developed this response to four specific areas of the Government's consultation: AI Interaction, Information about AI Systems, AI Incidents, and AI Agents.

ICTC's response is structured around the principle that transparency is most effective when it is designed for a defined recipient and purpose. A disclosure that is technically available but difficult to understand or act upon will do little to reduce harm, improve accountability, or strengthen Canadians' trust in AI. To achieve the Government of Canada's objectives of enhancing Canadians' understanding and trust in AI, ICTC has outlined a series of ten mutually reinforcing recommendations.

Canada should adopt a risk-based, role-based, and context-sensitive framework to govern AI transparency requirements. Requirements should increase with the potential consequences of an AI system and the degree of authority and responsibility delegated to AI systems. AI developers, deployers, operators, and platforms will therefore have different but complementary responsibilities. Canada should combine binding requirements for consequential uses with standards, guidance, procurement conditions, research, and capacity-building for end users to ensure that AI transparency is embedded in Canada's broader AI industry and ecosystem. The framework should also distinguish public disclosure from information provided to affected individuals, business customers, auditors, and regulators.

ICTC's position draws on its prior recommendations calling for predictable risk-based governance, proactive disclosure, regulatory interoperability, accredited assurance, regulatory sandboxes, investments in AI literacy, and practical support for small and medium-sized enterprises to adopt, implement, and govern AI. For AI agents, ICTC proposes a framework of Graduated Autonomy and Agency.

An AI agent should receive increased operational or economic authority only after demonstrating competence, reliability, security, and accountability in the environment in which it operates. Active agents should have verifiable identities, bounded permissions, proportionate activity records, traceable delegation, effective revocation, and an identifiable person or organization that remains accountable.

By adopting these recommendations and implementing them across a coordinated mix of policy instruments, Canada can foster increased trust and confidence in AI systems and AI agents while enabling business adoption, catalyzing productivity, and improving business competitiveness and innovation on a global scale.

Summary of Recommendations

1. **Adopt a risk-based, role-based, and context-sensitive framework.** Transparency obligations should reflect the AI system's function, its deployment setting, the degree of assigned authority, and the potential consequences on people and organizations.
2. **Create a meaningful right to be informed.** Canadians should receive clear notice when they interact directly with AI or when AI materially influences a consequential decision, service, or process affecting them.
3. **Standardize layered information about AI systems.** Canada should develop adaptable documentation that serves individuals, business adopters, auditors, researchers, and regulators without requiring indiscriminate public disclosure of sensitive information.
4. **Establish a coordinated AI incident reporting and learning framework.** Serious incidents should be reported through a common framework that complements existing sectoral regimes, protects confidentiality, and publishes anonymized lessons and disclosures to enhance trust.
5. **Implement Graduated Autonomy and Agency for AI agents.** Authority delegated to AI agents should expand or contract according to demonstrated performance, risk, context, and the consequences of the actions an AI agent is permitted to take.
6. **Develop the trust layer for the AI agent economy.** The widespread creation and adoption of AI agents is creating new terrains for Canada should lead on AI agent identity and credentials, progressive authorization, delegation protocols, activity evidence, assurance, interoperability, and revocation.
7. **Use standards, assurance, sandboxes, and procurement.** Legislation should be supported by technical standards, independent verification, supervised experimentation, and federal purchasing requirements to ensure that transparency requirements are proportionate, relevant, and agile enough to address the rapid development of AI.
8. **Provide proportionate pathways for SMEs to catalyze adoption.** Governance templates, technical assistance, shared tools and resources, phased requirements, and trusted advisory services should reduce compliance costs while supporting responsible AI adoption.
9. **Invest in AI literacy and governance capacity.** Workers, consumers, businesses, and public institutions need the skills to interpret disclosures and act on them. They also need the foundational AI literacy to understand how AI works, how their sensitive and private data is handled, and the measures they can take to protect their privacy while using AI safely and responsibly.
10. **Coordinate across jurisdictions and evaluate outcomes.** Federal measures should align with provincial, territorial, and international frameworks and be assessed for comprehension, contestability, risk reduction, and accountability. Canada should look to trusted partners to align regulatory frameworks to enhance interoperability and reduce barriers to businesses.

Principles for Effective AI Transparency

The Government of Canada recognizes that transparency is a foundation for informed decision-making and accountability. The policy objective should therefore be to make the role, limits, risks, and governance of AI visible, useful, and relevant to the people and organizations that need that information, while also preserving and respecting the variations in AI systems and deployment contexts.

To that end, ICTC proposes several guiding principles to help the Government of Canada develop AI transparency frameworks that are role- and context-specific.

Principle	Application
Proportionality	AI transparency requirements should be proportional to risks, contexts, and roles. Stronger requirements should apply where AI affects rights, safety, employment, finances, access to essential services, sensitive data, or critical systems.

Principle	Application
Responsibility by role	Each actor within the AI value chain should be accountable for the information they control. Developers explain the system; deployers explain its use; operators account for real-world decisions and outcomes.
Meaningful disclosure	Information should be timely, understandable, accessible, and linked to a decision, right, review process, or corrective action.
Layered access	Public summaries, user notices, business documentation, and confidential regulatory records should contain different levels of detail.
Accountability	Transparency should identify who is responsible and how an affected person can obtain review or redress.
Interoperability	Canadian measures should align with international standards and major markets while protecting Canadian rights and interests.
Adaptability	Standards and guidance should be updated as technologies, evidence, and deployment practices change.

These principles are consistent with ICTC’s previous recommendation that Canada establish predictable, proportionate AI governance supported by proactive disclosure, independent oversight, regulatory sandboxes, and international interoperability. They also reflect ICTC’s view that public trust and AI adoption are mutually reinforcing. Businesses are more likely to invest in transformational technologies when obligations are clear, vendor information is comparable and reliable, and responsibility across the value chain is defined.

AI Interaction

When disclosure is needed

Canadians should know when they are communicating with an AI system that could reasonably be mistaken for a person. They should also be informed when AI materially influences a consequential decision or process, even if they do not interact with the system directly. This includes decisions concerning employment, credit and access to financial services, insurance, healthcare, education, government benefits, legal services, and access to essential services.

A universal requirement to disclose every low-impact use of AI would create administrative costs and disclosure fatigue without providing a corresponding benefit to end users. Routine back-end automation should not require individual notice unless it materially affects a person, processes sensitive information, or introduces a significant risk. The trigger should be grounded in the nature and consequence of the use, rather than the mere presence of AI; for example, if an AI system has a material role in making a decision that impacts an individual, then disclosure should be mandated to ensure the individual has full awareness of the impact of AI.

The table below includes several examples of different types of human-AI interaction and ICTC’s recommended disclosure requirement. While non-exhaustive, the examples cited represent common use cases of human-AI interaction.

Type of interaction	Recommended requirement
Direct human-like communication, such as chatbots or voice AI customer service use cases.	Disclose at the start of the interaction that the user is communicating with AI. Keep the disclosure accessible throughout the interaction and provide the option to communicate with a human instead.
Consequential decision or process, such as pre-screening job applications	Explain that AI is being used, its role, the degree of human review, the principal information considered, and the route to contest or escalate the outcome.
Use of sensitive personal information	Explain how information is collected, used, retained, corrected, and removed, subject to applicable privacy law.
Low-impact back-end use	No general notice unless the use materially affects a person or creates a significant privacy, safety, rights, or security risk.

Responsibilities across the AI value chain

AI deployers carry primary responsibility for user-facing disclosure because they determine the purpose, context, audience, data practices, and degree of human oversight in each use case. AI developers should provide deployers with accurate information about capabilities, limitations, failure modes, and appropriate disclosure methods. Employers and other operators should also disclose consequential organizational uses. Platforms should support common technical and presentation standards where interactions occur through their infrastructure.

Defining meaningful disclosure

- The Government of Canada should consider defining meaningful disclosure to establish standards that apply across use cases, systems, and contexts. A meaningful disclosure should enable the recipient to adjust their behaviour, verify and correct information, seek human assistance, or challenge an outcome. A generic statement that an organization uses AI somewhere in its operations does not meet that standard. Features of a meaningful disclosure may include providing clear notice before or at the beginning of the interaction with AI, rather than burying it in terms of service.
- Using plain language and a format suited to the interface, including spoken notice for voice AI systems.
- Identifying the role of the AI system and whether it can make commitments or decisions on behalf of the organization.
- Explaining whether personal information is used and how it is protected.
- Describing the level of human oversight involved and providing a practical route to human review where warranted.
- Making notices accessible, available in both official languages, and tested with intended users.

Organizations should remain responsible for AI systems they deploy to communicate with the public. A business should not be able to avoid responsibility for representations made through its authorized chatbot solely because the system generated an erroneous response. Clear accountability will also create incentives to define and delimit the system's authority, monitor its performance, and provide timely human escalation.

Recommended actions

- Establish a baseline right to know when a person is communicating with an AI system that could reasonably be mistaken for a human.
- Require enhanced notice and human review mechanisms for consequential uses.
- Develop a layered disclosure model consisting of a short notice, a concise explanation of the system's role, access to detailed information, and an escalation route where appropriate.

- Issue sector-specific guidance and reusable templates, including simplified tools for SMEs.
- Define a meaningful disclosure standard to ensure consistency across use cases.

Information About AI Systems

Information should reflect the needs of the recipient and the use case

Canada should promote standardized, layered, and audience-specific information about AI systems. The objective should be to provide the relevant information needed by users to assess suitability, risk, and accountability, rather than to maximize the volume of information disclosed.

As noted in the Government of Canada's discussion paper, information about AI systems can be grouped into two broad considerations: information about what an AI system can and cannot do and information about how the AI system was designed, including the data on which it was trained. Information about both considerations should be made available to recipients to support AI transparency and build trust, supported by a common documentation framework to ensure consistency, accessibility, and relevance.

A common documentation framework

Standardized documentation would make systems easier to compare, reduce procurement costs, and help smaller organizations conduct due diligence. Templates should remain adaptable to different technologies and sectors. Canada should also consider developing shared documentation templates with trusted partners, such as the EU, to ensure interoperability and reduce barriers to trade.

The documentation framework can be tiered according to the relative impact of AI systems, use cases, and the needs of the target recipient. For higher-impact systems, documentation should be thorough and address the following areas:

- System identity, version, provider, and responsible organization
- Intended purposes, appropriate uses, and material restrictions
- Capabilities, performance, limitations, and known failure modes
- Performance across relevant populations and deployment conditions
- Data sources, provenance practices, and material representational gaps
- Treatment of personal, proprietary, copyright-protected, and Indigenous data
- Safety, security, privacy, and bias testing
- Human oversight, monitoring, and update requirements
- Incident reporting contacts and procedures
- Material changes following deployment
- Independent evaluation, certification, or assurance status

Training data and confidential information

Meaningful input transparency does not require the disclosure of every item in a training dataset. AI developers should, however, disclose the principal categories and sources of training data, relevant collection periods and geographic coverage, whether personal or copyright-protected information was included, material limitations in representation, and the processes available for individuals and rights holders to raise concerns. AI developers should also disclose if end-user data will be used to further train AI systems and models. Downstream organizations, such as AI deployers, should receive enough information to assess suitability and legal risk.

Canada should distinguish legitimate confidentiality from information required to protect rights, assess safety, or support regulatory oversight. Public documentation can remain high-level where necessary, while regulators and accredited auditors receive more detailed information under appropriate confidentiality and security protections.

Recommended actions

- Develop standardized and adaptable model and system documentation templates aligned with international standards.
- Set stronger documentation requirements for high-impact systems and systems that process sensitive information.
- Establish minimum information rights for businesses procuring third-party AI systems.
- Require notice when a material update changes a system's risk, performance, or appropriate use.
- Develop accreditation and independent assurance pathways for higher-risk applications.
- Use federal procurement to require documentation, testing evidence, update notices, and contractual clarity.
- Provide shared due diligence tools, procurement guidance, and trusted advisory support for SMEs and non-profit organizations.

Documentation by itself is not enough; it must be paired with organizational capacity building to ensure businesses, individuals, and organizations are equipped to interpret documentation and to understand how it applies to their intended usage. Businesses need staff who can evaluate vendor claims, integrate systems responsibly, monitor performance, and respond to incidents. Integrating sector-specific AI training, SME adoption support, and stronger AI literacy across the workforce can support the usability of documentation.

AI Incidents

A coordinated reporting and learning framework

Canada does not yet have a systematic view of AI incidents and emerging risk patterns. Existing incident reporting regimes used for medical devices, transportation, consumer products, privacy, and cybersecurity capture some events, but they use different definitions and were not designed to provide a cross-sector understanding of AI risk.

Canada should establish a coordinated framework that complements these regimes and avoids duplicate reporting. The framework should support accountability for serious incidents and collective learning from lower-level events. Mandatory reporting is appropriate when harm is serious or imminent. Protected voluntary reporting can encourage organizations to share near misses and contained failures that would otherwise remain invisible, while also improving ecosystem-wide understanding and trust.

Definition of a serious AI incident

A serious AI incident or near-miss should be defined by its consequences and the potential consequences. A reportable event should include an AI-related malfunction, misuse, malicious act, security compromise, or governance failure that causes, materially contributes to, or creates an imminent risk of:

- Death or serious harm to health or safety
- A significant infringement of human or constitutional rights
- Material discrimination
- Significant financial loss or fraud
- Unauthorized access to or disclosure of sensitive information
- Compromise of critical infrastructure or essential services
- Loss of meaningful human control over a consequential system
- Serious environmental or property damage
- Large-scale manipulation or disruption of democratic and civic processes
- Systemic harm affecting a substantial population
- Unauthorized exercise or delegation of an AI agent's authority

Reporting of AI incidents and near-misses should be tiered according to the severity, likelihood, and scale of the consequences. For example, Canada may adopt a tiered system that ranks AI incidents from Critical, Serious, Near Miss, and Operational Trend. The table below illustrates how these tiers may be differentiated and the corresponding reporting requirement.

Tier	Illustrative event	Reporting approach
Critical	Death, severe injury, critical infrastructure disruption, major rights violation, or significant security compromise	Rapid mandatory confidential report to the appropriate authority
Serious	Material financial harm, discriminatory outcome, major privacy breach, or consequential unauthorized action	Mandatory report within a defined period
Near miss	Contained failure, attempted exploit, or error detected before harm	No-fault reporting
Operational trend	Recurring errors, complaint patterns, or reliability changes	Periodic aggregated reporting

Information and responsibility

The Government of Canada should also establish standards governing the information disclosed in reports of AI incidents, according to the severity of the incident. Reports should identify the system and version, developer, deployer, operator, deployment context, affected parties, timing, scope, nature of the harm, oversight arrangements, containment measures, root cause when known, and corrective action. Security-sensitive information should be disclosed through coordinated processes that do not expose an unremediated vulnerability.

Reporting responsibility should follow knowledge and control. Deployers and operators are often best placed to identify incidents in real-world use. Developers are better placed to identify model-level failures or vulnerabilities that may affect multiple deployments. Both should report when each holds material information, and contracts should require timely information-sharing across the value chain.

Institutional design

Canada should establish a federal coordinating function or trusted reporting agency that routes cases to the appropriate regulator and works with the Canadian AI Safety Institute, privacy authorities, and the Canadian Centre for Cyber Security. The framework should use a common taxonomy and publish anonymized trend reports, safety alerts, and lessons learned. A multi-stakeholder advisory body could help maintain definitions, thresholds, and reporting standards as evidence develops.

Appropriate confidentiality and safe harbour protections should encourage good-faith reporting without shielding negligence, concealment, or misconduct. Canada should also align its taxonomy and reporting formats with trusted international systems to reduce duplicate compliance and improve cross-border learning, while also embedding confidentiality measures on a limited basis to reduce the risk of broadcasting systemic vulnerabilities to hostile actors.

Recommended actions

- Map existing federal, provincial, and territorial reporting regimes before introducing new requirements.
- Develop a common Canadian AI incident taxonomy and harmonized reporting form.

- Require reporting of serious incidents and support protected reporting of near misses.
- Clarify confidentiality, liability, and good-faith reporting protections.
- Require corrective action and notification where a vulnerability may affect other deployments.
- Publish anonymized findings and risk trends for industry, researchers, and the public.
- Pilot reporting processes in regulatory sandboxes and selected high-impact sectors.
- Provide simplified tools and guidance for SMEs.

AI Agents

Graduated Autonomy and Agency

AI agents require a distinct approach within Canada's AI transparency frameworks because they can take actions, use tools, access data, transact, and delegate tasks with greater autonomy than other AI systems. The relevant risk depends on more than the sophistication of the underlying model. It arises from the authority granted to the agent, the assets and systems it can access, the degree of independent action permitted, and the consequences of error or misuse.

ICTC proposes Graduated Autonomy and Agency as the organizing principle for AI agent governance. Greater operational or economic authority should follow from demonstrated competence, reliability, security, and accountability within the agent's operating environment.

Illustrative stages of authority

Stage	Illustrative authority	Required controls
Observe	Retrieve or summarize information	Basic identity, privacy, security, and access controls
Recommend	Propose an action	Disclosure, source information, and human review
Act with confirmation	Prepare a transaction that requires approval	Strong authentication, clear confirmation, and activity evidence
Limited delegated action	Purchase below a threshold from approved vendors	Verified identity, scoped credentials, limits, monitoring, and revocation
Expanded organizational action	Execute larger transactions or modify approved systems	Independent assurance, continuous monitoring, segregation of duties, and incident reporting
Multi-agent delegation	Delegate a defined task to another approved agent	Chain-of-agency records, delegation limits, credential rules, and revocation propagation

Progression towards subsequent stages should not be automatic; rather, authority should expand, contract, expire, or be revoked when context, performance, system behaviour, credentials, or risk changes. Furthermore, a material model update should trigger a reassessment when it may alter the agent's behaviour or reliability.

The trust framework for economically active agents

Verifiable identity: An active agent should have a verifiable identity that establishes what the agent is, its provider and version, the person or organization on whose behalf it operates, and the status of its credentials and authority. Requirements should be strongest where agents have direct control over

sensitive tasks, data, or information, such as accessing sensitive information, entering agreements, modifying critical systems, or interacting with other agents.

Progressive authority: Permissions should be explicit, purpose-limited, and time-limited where appropriate, and constrained by transaction value, data type, vendor, system, or operating environment. Authority should be reviewable and capable of expansion, contraction, suspension, and expiry. An agent's technical ability to perform an action should not be treated as authorization to perform it.

Activity evidence: Organizations should retain enough evidence to develop a history of activity and establish which agent acted, on whose behalf, under which permission, what consequential action it took, when it acted, which system or agent it engaged, whether approval was required and obtained, and what result followed. This should not require indiscriminate retention of prompts, personal information, proprietary information, or internal model reasoning. The objective is traceability of consequential actions for ongoing monitoring, improvement, and reporting when incidents occur.

Chain of agency: When one agent delegates to another, records should identify the original accountable person or organization, the originating agent, the authority initially granted, each subsequent delegation, any narrowing or expansion of purpose or duration, and the agent that performed the final action. A delegated agent should never receive broader authority than the delegating agent possesses.

Revocation: Credentials and permissions must be capable of being restricted, suspended, expired, or revoked. Revocation should propagate through delegated chains and should be available when behaviour changes, credentials expire, a system is materially updated, a security compromise occurs, or an incident reveals that existing authority is unsafe.

Human and organizational accountability: Every consequentially active agent should be linked to an identifiable accountable person or organization. The accountable party should define the agent's authority, monitor performance, respond to incidents, provide redress, manage delegation, and revoke access when necessary. AI developers should remain responsible for defects, misleading representations, and failures within their control, while AI deployers should remain responsible for deployment choices, permissions, monitoring, and organizational use. Importantly, increasing the autonomy of an AI agent should not make human accountability disappear.

A risk model based on authority and consequence

Oversight should be calibrated using four factors:

- **Authority:** the actions the agent is permitted to take
- **Assets:** the money, data, systems, infrastructure, or rights it can access
- **Autonomy:** whether it can act without contemporaneous human confirmation
- **Consequences:** the harm that could result from error, compromise, or misuse

An agent that summarizes documents should face lighter requirements than an agent authorized to spend a limited amount with approved vendors; the latter should have verified identity, transaction records, defined limits, and straightforward revocation. An agent that can transfer substantial funds, access health information, enter contracts, modify critical systems, or delegate authority should be subject to stronger assurance, monitoring, approval, and incident-reporting requirements with built-in human oversight at regular gates.

Canada's opportunity in the agent trust layer

The development of AI agents and their growing usage requires the development of cohesive and joint regulatory standards and frameworks. Canada can seek to lead in the development of the regulatory infrastructure that allows organizations and markets to trust AI agents. This may include agent identity and credentials, progressive authorization, delegation protocols, activity records, independent assurance,

interoperability, and revocation. These capabilities will be necessary whether the underlying models are developed in Canada or abroad.

A Canadian agent trust layer would align with Canada's objectives with regard to digital sovereignty, cybersecurity, commercialization, responsible governance, and participation in international standards. It could also create opportunities for Canadian firms in digital identity, privacy-enhancing technology, payments, cybersecurity, governance software, observability, testing, certification, insurance, and secure cloud infrastructure.

Recommended actions

- Develop a national trust framework for economically and operationally active AI agents.
- Establish interoperable standards for agent identity and credentials.
- Define progressive authorization and delegation protocols.
- Require proportionate evidence of consequential actions and chains of agency.
- Require effective credential expiry, suspension, and revocation.
- Clarify that deploying organizations remain accountable for authority granted to agents.
- Develop liability and redress principles for unauthorized or erroneous agent actions.
- Use regulatory sandboxes to test agentic commerce, enterprise delegation, and machine-to-machine transactions.
- Promote independent assurance for agents with significant economic or operational authority.
- Use federal procurement as an early proving ground for agent identity, permission controls, activity evidence, and revocation.
- Support SMEs with shared identity, assurance, and compliance infrastructure.

Implementation

Use a coordinated mix of policy instruments

No single instrument will address the range of transparency needs identified in the consultation. Canada should use legislation where enforceable rights, serious harms, or mandatory reporting are involved. Standards and guidance can define technical and operational practices. Procurement can create early demand for responsible products. Sandboxes can generate evidence where technology and standards remain immature. Literacy and advisory programs can help users and organizations act on the information they receive.

Instrument	Best use
Legislation and regulation	Baseline rights, consequential uses, serious incident reporting, accountability, and redress
Standards and guidance	Documentation formats, identity, credentials, logging, delegation, assurance, and interoperability
Regulatory sandboxes	Testing emerging requirements for agents and other rapidly changing applications
Public procurement	Creating demand for documentation, human oversight, incident readiness, and agent controls
Research and evaluation	Testing comprehension, effectiveness, unintended effects, and technical feasibility
Literacy and technical assistance	Helping individuals and organizations interpret disclosures and implement responsible practices

Support small and medium-sized enterprises

SMEs should not be exempt from protecting people against serious harm, but compliance pathways should reflect their capacity and their role in the value chain. A small business deploying a third-party system should not be expected to produce technical information available only to the developer. Conversely, vendors should not shift responsibility for model-level information to customers that cannot reasonably obtain it.

The Government should support SMEs through:

- Standard notices, documentation templates, and contractual clauses
- Sector-specific implementation guides and adoption diagnostics
- Trusted advisory services and technical assistance
- Phased implementation and proportionate recordkeeping
- Shared testing, assurance, identity, and incident-reporting infrastructure
- Regulatory sandboxes and supervised pilot programs
- Training for procurement, governance, cybersecurity, and incident response

Coordinate across jurisdictions

Federal measures should use common definitions and templates wherever possible and should be developed with provinces, territories, sectoral regulators, standards bodies, Indigenous partners, industry, labour, researchers, and civil society. International alignment is especially important for documentation, incident taxonomies, agent credentials, assurance, and cross-border delegation. Canada should seek interoperability with trusted partners while retaining the ability to protect Canadian rights, data, infrastructure, and economic interests.

Measure whether transparency works

Compliance should not be measured only by whether a notice or document exists. The Government should assess whether intended recipients understand the information, whether it improves procurement or user decisions, whether affected people can obtain review and redress, whether incident reporting produces preventive action, and whether agent controls prevent unauthorized activity. Requirements should be revised when evidence shows that they create fatigue, confusion, unnecessary cost, or ineffective compliance.

Conclusion

Meaningful AI transparency can strengthen public confidence, protect Canadians, encourage responsible AI usage, and give businesses greater certainty to adopt AI responsibly. ICTC recommends a risk-based, role-based, and context-sensitive framework that makes information understandable and actionable for its intended recipients, with clear responsibilities across the AI value chain. This framework should establish a meaningful right to be informed about direct AI interactions and consequential uses, standardize layered information about AI systems, and coordinate serious incident reporting and protected reporting of near misses to support accountability and collective learning. For AI agents, Graduated Autonomy and Agency should link delegated authority to demonstrated competence, reliability, security, and accountability. Verifiable identities, bounded permissions, traceable actions and delegation, and effective revocation should underpin a Canadian agent trust framework, with an identifiable person or organization remaining accountable.

Implementation will require coordinated legislation, technical standards, independent assurance, regulatory sandboxes, and public procurement, supported by investments in developing Canadians' AI literacy and governance capacity. Practical tools, shared infrastructure, and proportionate compliance pathways for SMEs will help translate these requirements into responsible adoption that supports business productivity, innovation, and competitiveness without compromising security or accountability. Coordination across Canadian jurisdictions and alignment with trusted international partners can improve interoperability while protecting Canadian rights and interests.

The success of Canada's AI Transparency framework should be measured by whether it improves understanding and decision-making, enables review and redress, and reduces harm. Together, these measures can build trust in AI while supporting productivity, innovation, digital sovereignty, and Canada's competitiveness in the global economy.

Referenced ICTC Publications

Anani, Namir, and Erik Henningsmoen. *A Roadmap for Canada's Digital Economy to 2030*. Ottawa: Information and Communications Technology Council, April 2025. <https://ictc-ctic.ca/reports/a-roadmap-for-canadas-digital-economy-to-2030>.

Garmendia, Amaya, Todd Legere, and Noah Lubendo. *AI Sovereignty and Economic Growth: Strengthening Transatlantic Leadership Between the EU and Canada*. (AI, Data and Robotics Association (Adra) and Information and Communications Technology Council, May 2025). <https://ictc-ctic.ca/reports/ai-sovereignty-and-economic-growth>.

Henningsmoen, Erik, Noah Lubendo, Mairead Matthews, and Heather McGeer. *Advancing Canada's National Sovereign AI Compute Capacity: An ICTC Policy Brief*. (Ottawa: Information and Communications Technology Council, July 2025). <https://ictc-ctic.ca/policy-briefs/advancing-canadas-national-sovereign-ai-compute-capacity>.

Information and Communications Technology Council. *From AI for All to AI in Action: Implementing Canada's National AI Strategy for Shared Prosperity*. (Ottawa: Information and Communications Technology Council, June 2026). <https://ictc-ctic.ca/reports/from-ai-for-all-to-ai-in-action>.

Information and Communications Technology Council. *Seven Key Recommendations for Canada's Renewed National AI Strategy*. (Ottawa: Information and Communications Technology Council, October 2025). <https://ictc-ctic.ca/reports/seven-key-recommendations-for-canadas-renewed-ai-strategy>.