



Submission to the consultation on advancing AI transparency in Canada

Global Privacy Rights 23 September 2026

Summary

The discussion paper treats transparency as information that developers and deployers provide. That is necessary, and it is not sufficient. Once a disclosure has been made, the individual holds nothing that shows what they were told, by whom, or when. When a dispute arises, the only account of the disclosure is the account of the organisation that made it.

We recommend that Canada treat transparency as evidence as well as information: the accountable organisation identified before any interaction, every disclosure versioned so that it can be retrieved as it stood on a given day, and a receipt of the disclosure available to the individual. These measures apply equally to AI interactions, to information about AI systems, and to the activity of AI agents, and they can be built on standards that already exist.

1 The gap: evidence held by the individual

Bill C-36, the Protecting Privacy and Consumer Data Act, which the paper cites, shows the gap clearly. Our published assessment of the Bill at first reading (WHISSPR Report: Bill C-36 Assessment, Global Privacy Rights, 2026, <https://globalprivacyrights.org/reports/whisspr-bill-c-36-assessment-2026.pdf>) found the following.

- Section 15(3) requires five items to be provided before consent is sought. Nothing in the Bill requires that provision to be recorded, versioned or evidenced.
- Section 8(1) makes the contact information of the designated individual available only on request. No provision requires the identity of the accountable organisation to be given before collection.
- The word receipt does not appear in the Bill. Every record it creates is held by the organisation and disclosed to the regulator on request.

Quebec's Act respecting the protection of personal information in the private sector, as amended by Law 25 and in force, shows the same asymmetry in a statute that already applies. Our published assessment of that Act (WHISSPR Report: Quebec Law 25

Assessment, Global Privacy Rights, 2026, <https://globalprivacyrights.org/reports/whisspr-quebec-law-25-assessment-2026.pdf>) found that the word receipt appears four times and every occurrence concerns the receipt of a request or of a decision. Section 54 requires the Commission d'accès à l'information to send its decision to the parties by any means providing proof of the date of receipt. The Act requires proof of receipt when the regulator communicates with a party, and requires none when an enterprise gives notice to an individual.

The result is that an organisation that complies fully and one that does not comply at all produce the same evidence for the individual, which is none. AI transparency measures built on the same model will have the same weakness: a disclosure that an individual is interacting with an AI system leaves no trace that the individual can later rely on.




Figure 1. What Bill C-36 records, and who may read it. Every record the Act creates is held by the organisation and disclosed to the Commission on request. The individual holds nothing: no receipt, no notice record, no version, no event log.

2 AI interaction

When is it most important to know, and what makes a disclosure meaningful? A disclosure is meaningful when it can be checked. That requires four properties:

1. the accountable organisation is identified, in a published record that can be retrieved before the interaction begins, and not only on request;
2. the disclosure is versioned, with a publication time and an integrity value, so that the version in effect on a given day can be retrieved later;
3. the individual can obtain a receipt of the disclosure, generated when it is presented, that references that version;
4. the receipt can be obtained and verified without creating an account or identifying oneself.

Responsibility should sit with the deployer that presents the interaction, because that is the party the individual meets. Information from the developer should be referenced in the deployer's disclosure, not left for the individual to find.




Figure 2. Where Bill C-36 places notice in the identification sequence. No controller identity is required at any point, section 15(3) applies only where consent is required, and section 62 sits outside the sequence because it carries no stated timing.

3 Information about AI systems

Disclosures about capabilities, limitations and development change as systems are updated. A disclosure that can be silently replaced cannot support an informed choice, a complaint or an audit. Each published version should carry its publication time and an integrity value, and superseded versions should remain retrievable for as long as any interaction relied on them.

This addresses the tension the paper raises with confidential business information. What must be evidenced is the disclosure the organisation chose to make, not the underlying trade secret. Versioning a public statement exposes nothing that was not already published.

4 AI agents

On what should be recorded about agent activity, and disclosed to whom:

1. each agent should carry an identifier and a credential scoped to its task, which the service receiving the request can verify;
2. the authorisation the agent relies on should be recorded: who gave it, for which purpose, and for how long;
3. approvals given by a person at critical points should be recorded in a form that cannot be altered afterwards;

4. each of these records should reference the notice the individual was given, so that the authorisation can be traced to the disclosure it depends on;
5. the individual on whose behalf, or about whom, an agent acts should be able to obtain a record of the authorisation relied upon, not only the operator and the regulator.

Other jurisdictions are moving here. China's AI Safety Governance Framework 3.0, published by the National Technical Committee 260 on Cybersecurity in September 2026, calls in Appendix 2 for a unique identity and scoped credentials for each agent, and for tamper-proof records of human approval. Those records are held by the operator. None is held by the individual. Canada has the opportunity to set the stronger model.

The same holds across borders. Where an agent's activity, or the data it relies on, moves to another jurisdiction, the individual should be able to learn to whom, to which jurisdiction, and under what instrument. Bill C-36 reports only whether a transfer occurred, and the one route by which an individual might learn more is closed by section 70(6). Figure 3 sets out that asymmetry.

Activity records should not themselves become a new source of surveillance. Logs of agent activity should exclude identifying data about individuals that the purpose does not require, including network, device and session identifiers. The Blinding Identity Taxonomy, a Kantara Initiative Report, lists the categories concerned.



Four routes out. One route back, and it is gagged.

Section 62(2)(d) tells the individual whether or not a transfer occurred. Not to whom, not to which jurisdiction, and not under what instrument.

Source: Bill C-36, ss.43, 44, 62(2)(d), 63, 70(6), 82, 133; Aeronautics Act s.4.83(1) as amended by cl.27.
Assessment: Transparency Lab, published by Global Privacy Rights.

Figure 3. Transborder asymmetry under Bill C-36. Four routes carry information to a foreign authority. One route carries anything back to the individual, and section 70(6) gags it. Section 62(2)(d) reports whether a transfer occurred, not to whom, not to which jurisdiction, and not under what instrument.

5 Operational transparency: the security of the observation

Security requirements for AI systems ask whether unauthorised access occurs. They rarely ask whether authorised access is accountable. A system can satisfy every control in a security framework while access to what it records is standing rather than granted, logged only to the party making the request or not logged at all, unbounded in time and scope, and invisible to the person the records are about.

That distinction matters more for AI systems than for the systems that preceded them, because the records an AI system generates are not limited to what a person provided. They include what was inferred, what was retained as context, and what was disclosed onward to a model or a tool. Who may read those records, on what threshold, and with what trace, is a transparency question before it is a security one.

Four properties make an access accountable, and each can be stated as a requirement. The request carries a credential at the point of access rather than a standing permission. The access is logged to a party that is not the requester. It is approved actively rather than by prior arrangement. And it is bounded in time and in scope. A system with all four can be inspected after the event. A system with none can be entirely secure and entirely unaccountable at the same time.

6 Rights that scale with the system

Transparency obligations are usually set at one level for every system, which makes them too heavy for a simple tool and too light for a system that infers, profiles, or moves data across borders. Deriving the rights controls from the transparency record itself resolves that. Where the disclosure is published in a machine-readable form, the controls available to the individual can be expressed in the same form and attached to it, so that the more inference, profiling or cross-border transfer a system performs, the more it must disclose and the stronger the controls that attach.


Global Privacy Rights has developed this as derivative rights controls, derived from an open code of practice profile grounded in international law and existing standards. Two consequences follow. Rights stop being entitlements an individual pursues after the fact and become controls reachable from the notice at the time it is given. And a supervisory authority gains something it does not have today, a standing and inspectable statement of who is processing and under what authority, comparable across organisations rather than examined one complaint at a time.

7 Coherence with provincial and international frameworks

1. Quebec's Law 25 requires the title and contact information of the person in charge of the protection of personal information to be published on the enterprise's website (section 3.1). It is the strongest such requirement in North America and it is already in force. A federal measure should meet that standard, not fall below it. Assessed in the report cited in section 1, and compared with Bill C-36 and with

China's AI Safety Governance Framework 3.0 in the WHiSSPR Comparative: Québec, Canada and China, Global Privacy Rights, 2026, <https://globalprivacyrights.org/reports/whisspr-comparative-quebec-canada-china-2026.pdf>

2. The European Union's General Data Protection Regulation, Articles 13 and 14, requires the identity and contact details of the controller when personal data is collected. Canadian measures that identify the accountable organisation before collection support interoperability with that framework.
3. The paper does not refer to the Council of Europe's instruments, including Convention 108+ and the Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law. Both bear directly on transparency and should inform a Canadian approach.
4. Existing standards already specify much of what is needed: ISO/IEC TS 27560:2023, Privacy technologies, consent record information structure; and ISO/IEC 29184:2020, Online privacy notices and consent. Building on them avoids a Canadian format that does not interoperate.
5. Bill C-36, sections 92 to 96, creates a route for approved codes of practice and certification programs, including independent verification of compliance (section 93(1)(d)) and public approval decisions (section 95). An AI transparency code of practice built on the measures above, and on the standards in item 4, could be approved through that route without waiting for new legislation.
6. The Government of Canada already publishes an AI Register for federal institutions, launched on 28 November 2025 as a minimum viable product on the Open Government Portal. The discussion paper does not refer to it. The April 2026 data lists 412 systems from 43 institutions. Of the 88 systems recorded as involving personal information, 31 name a personal information bank. Of the 54 used by or with the public, 22 are recorded as notifying the public. A register describes systems; it does not give an individual a record of an interaction with one. The register model should be extended to AI systems offered by private-sector organisations, with a published record of the accountable organisation, versioned entries, and a link from each entry to the notice and receipt the individual receives.



Source: Bill C-36, ss.92 to 96. Assessment: Transparency Lab, published by Global Privacy Rights.

Figure 4. The route opened by Bill C-36, sections 92 to 96. Any entity may apply, including one not subject to the Act. A code may add controller identification before collection, a versioned notice record, and a receipt held by the individual, with independent verification written into section 93(1)(d).

8 Small and medium-sized enterprises

A published organisation record, a versioned notice and a receipt are small, standard records. They do not require an organisation to operate identity systems or to hold more data about individuals. Open, standard formats keep the cost low and allow small enterprises to use shared tools.

The position these recommendations address

These recommendations address the position of the individual. A collective authority position exists alongside it, held by a community over its own data, knowledge, language and territorial information, and it is not reducible to the individual rights of that community's members. It is not addressed here, and it should be defined by the rights holders concerned rather than on their behalf.

Recommendations

1. Require the organisation accountable for an AI interaction to be identified, in a published and retrievable record, before the interaction begins.
2. Require AI transparency disclosures to be versioned, with a publication time and an integrity value, and superseded versions to remain retrievable.
3. Give individuals evidence: a receipt of the disclosure, available without an account or identification. Specify the record so that the accountable operator, the purpose,

and the authority relied on are carried as separate fields, so that an authority other than the individual can be named in the same structure rather than requiring a second one.

4. Derive the rights available to an individual from the published transparency record and scale them to the processing, so that systems performing more inference, profiling or cross-border transfer carry stronger controls.
5. For AI agents, require scoped, verifiable credentials and unalterable approval records, linked to the notice the individual was given, and make the authorisation record available to the individual.
6. Require access to the records an AI system generates to be credentialed at the point of request, logged to a party that is not the requester, actively approved rather than standing, and bounded in time and scope.
7. Use the codes of practice and certification route in Bill C-36, sections 92 to 96, to approve an AI transparency code of practice built on ISO/IEC TS 27560:2023 and ISO/IEC 29184:2020, with independent verification.
8. Align these measures with Quebec's Law 25 and take account of Convention 108+ and the Council of Europe Framework Convention on Artificial Intelligence.
9. Extend the federal AI Register model to AI systems offered by private-sector organisations, with a published accountability record, versioned entries, and a link to the notice and receipt given to the individual.

Global Privacy Rights would welcome the opportunity to discuss these proposals, and can demonstrate the measures recommended here in working form: an accountable organisation identified in a published record before an interaction begins, a versioned disclosure retrievable as it stood on a given day, and a receipt obtainable without an account or identification.

Mark Lizar Global Privacy Rights info@globalprivacyrights.org

About Global Privacy Rights

Global Privacy Rights is a Canadian not-for-profit. It conducts and publishes research on what services actually disclose, through Transparency Lab, its lab, with Global Privacy Rights as publisher of record. Published assessments are available at <https://globalprivacyrights.org/research> and cover Bill C-36, Quebec's Law 25, and China's AI Safety Governance Framework 3.0, with a comparative of all three produced on one method.

In Montreal, Global Privacy Rights is the operator and developer of the Operational Transparency for Security and Privacy of AI signalling registry, built as public infrastructure and running as a research and development pilot. It is the surface on which the measures recommended here are tested in the open, using open public standards throughout. It aims to conform with the guidance of ISO/IEC 27091, Cybersecurity and privacy, Artificial intelligence, Privacy protection, which is at final draft stage and not yet published, and with ISO/IEC 42001 for artificial intelligence management system lifecycle governance. Level 0

of its assurance ladder (<https://globalprivacyrights.org/assurance>) is live, where a listing is self-asserted and says so; verification begins at Level 2 and is not yet built.

It publishes its own controller identification record at <https://globalprivacyrights.org/.well-known/transparency/cir.json>, and states the assurance level of its own listings rather than claiming one, on the same terms it asks of any organisation it assesses.

Figures and attribution

Figures are drawn from the WHISSPR Report: Bill C-36 Assessment, Global Privacy Rights, 2026, available at <https://globalprivacyrights.org/reports/whisspr-bill-c-36-assessment-2026.pdf>. The assessment was conducted by Transparency Lab, the Global Privacy Rights lab, and published by Global Privacy Rights as publisher of record. Clause references are to Bill C-36 as at first reading.

The assessments this submission draws on are published in full by Global Privacy Rights at <https://globalprivacyrights.org/research>: the Bill C-36 assessment, the Quebec Law 25 assessment, an assessment of China's AI Safety Governance Framework 3.0, and a comparative of the three. Each was produced on the same method and each records its own sources and limits.

This submission and its figures may be reproduced with attribution.

© 2026 Global Privacy Rights.