

Transparency and Disclosure Rules for Government Use of Artificial Intelligence

February- 2026

DEMOCRATIC GOVERNANCE | PRIVACY & DATA RIGHTS | TECHNOLOGY GOVERNANCE

Civic Sketches is a series of draft policy proposals by Lawrence Nault — an independent researcher and policy writer based in Alberta's Badlands. These proposals are working drafts, not party platforms or final answers. They are intended to invite discussion, challenge, and refinement across all levels of government and the broader public.

This policy proposal is published under the Creative Commons Attribution 4.0 International Licence (CC BY 4.0). You are free to copy, share, adapt, translate, and build upon this work for any purpose, including use by governments, organizations, advocates, researchers, and members of the public, provided appropriate credit is given to Lawrence Nault and any changes are clearly identified. Full licence: creativecommons.org/licenses/by/4.0/

The core principle

If artificial intelligence helps shape public law, public policy, public analysis, public recommendations, or public decisions, then the model identity, prompts, outputs, and reliance trail should be treated as public records, subject only to narrow existing exceptions.

The proposed reforms

1. Define AI interaction records as public records

Alberta should amend its access-to-information law to expressly include AI-related records when AI is used for public business.

This should include:

- prompts;
- instructions;
- system messages;
- templates;
- outputs;
- summaries;
- draft text;
- recommendations;

- classifications;
- model identity;
- vendor information;
- system version;
- records showing human review, editing, approval, rejection, or reliance.

2. Create a duty to preserve AI records

Public bodies should be required to retain enough information to explain material AI use.

At minimum, records should include:

- the date of use;
- the purpose of use;
- the system or model used;
- the prompt or instruction set;
- the output received;
- any major edits made by officials;
- whether the output was relied on;
- any privacy, security, fairness, or validation review.

A public body should not be able to avoid public access simply by failing to keep the records.

3. Establish a public AI registry

Alberta should create a searchable public registry of significant AI systems used by public bodies.

The registry should identify:

- the public body using the system;
- the vendor or model name;
- the general purpose of use;
- whether the system is external or internally hosted;
- whether personal or confidential information is involved;
- whether the use is high-impact;
- whether a privacy or fairness assessment was completed;
- whether human review is required.

This would allow citizens, journalists, legislators, researchers, and oversight bodies to see where AI is being used in public administration.

4. Require proactive disclosure for high-impact AI use

For high-impact uses, government should not wait for someone to file an access request.

Proactive disclosure should be required when AI is used to support:

- legislation;
- policy development;
- electoral or boundary analysis;

- benefits administration;
- enforcement priorities;
- licensing;
- inspection;
- eligibility decisions;
- other functions that affect rights, services, obligations, representation, or public entitlements.

The disclosure does not need to reveal protected private or security-sensitive information. But it should provide a plain-language summary of how AI was used, what tool was used, whether human review occurred, and how further records can be requested.

5. Reform procurement rules

All AI vendor contracts used by Alberta public bodies should protect public access.

Contracts should require:

- retention of prompt and output records;
- government rights to inspect, copy, export, and disclose records where required by law;
- clear rules on whether government information can be used to train vendor systems;
- logging of model or system changes;
- audit rights for high-impact uses;
- no contract terms that override lawful access rights.

Government should not be able to lose public records by outsourcing part of its reasoning process.

6. Keep existing exemptions narrow

This proposal does not eliminate existing protections.

Personal privacy, solicitor-client privilege, cabinet confidentiality, security-sensitive information, and genuine trade secrets should still be protected where the law requires it.

But there should be no blanket exemption simply because a record involves AI.

AI-related records should be reviewed like other records: carefully, specifically, and with disclosure as the default unless a lawful exemption applies.

7. Create an express review right

Applicants should have a clear right to seek review by Alberta's Information and Privacy Commissioner when a public body:

- refuses to confirm whether AI was used;
 - claims prompts or outputs are not records;
 - fails to preserve relevant AI records;
 - over-redacts AI-related materials using broad vendor-confidentiality claims.
- Oversight matters. Access rights are only meaningful if refusals can be challenged.

Why this matters

This proposal would not ban government use of AI.

It would not stop useful tools from being adopted.

It would not require reckless disclosure of private, privileged, cabinet-confidential, or security-sensitive information.

It would simply make sure that when AI is used inside public administration, democratic accountability keeps pace.

The risk is not only that AI will make mistakes. The risk is that AI-assisted government work may become impossible to examine after the fact.

If the public cannot know what system was used, what it was asked, what it produced, and how officials relied on it, then part of the governing process has moved outside meaningful scrutiny.

That should not be acceptable in a democracy.

Questions for discussion

Should every material government use of AI be logged?

Should high-impact AI use require proactive public disclosure?

Should citizens have a right to know whether AI helped shape a decision that affected them?

Should vendor contracts be allowed to limit public access to AI-related government records?

What information should appear in a public AI registry?

How should the law balance transparency with privacy, security, cabinet confidentiality, and legitimate trade secrets?

What resources would Alberta's Information and Privacy Commissioner need to oversee this properly?

Plain-language summary

If government uses AI to help shape laws, policies, analysis, recommendations, or decisions, the public should have a right to know how.

That means preserving the model name, prompts, outputs, and records showing whether officials relied on the result.

AI should not become a way for public decision-making to disappear into a black box.

The Policy: Transparency and Disclosure Rules for Government Use of Artificial Intelligence

Purpose: To modernize Alberta's access-to-information framework so that the use of artificial intelligence in public decision-making, policy development, legislative drafting, analysis, and recommendations is subject to meaningful transparency, record-keeping, and public access.

Executive Summary

Alberta's **Access to Information Act (ATIA)** applies to records in the custody or under the control of a public body, including records created by outside consultants or contractors where the public body has relied on them or has contractual rights to inspect or copy them. Alberta's own guidance states that records created by an outside contracted consultant may fall under a public body's control, especially where the public body has relied on them to a substantial extent.

This principle should be expressly extended and clarified for artificial intelligence systems used by public bodies. When a ministry, agency, or other public body uses an external or third-party AI model to produce analysis, recommendations, draft text, summaries, classifications, or decision-support outputs, the model is functioning as part of the government's reasoning process. In such cases, the identity of the AI system, the instructions or prompts used to direct it, the outputs it generated, and the records showing how public officials relied on those outputs should be treated as presumptively accessible public records, subject only to narrow and existing statutory exceptions.

The Office of the Information and Privacy Commissioner of Alberta has already recommended that Alberta consider public-sector transparency obligations for AI and specifically consider a publicly accessible, searchable AI registry. The Commissioner also emphasized the overlap between AI, administrative fairness, and good governance, particularly where AI supports public decision-making or policy evaluation (Office of the Information and Privacy Commissioner of Alberta [OIPC], 2025).

This proposal recommends targeted reforms to Alberta's ATIA, related regulations, procurement rules, and records-management practices to ensure that public bodies cannot avoid accountability by outsourcing part of their analysis or drafting to opaque external systems.

Problem Statement

Alberta's current access framework was not written with contemporary generative or decision-support AI systems in mind. The ATIA already reaches records under a public body's control, and Alberta's guidance recognizes that contractor-created records can fall within that concept. However, the law does not yet clearly address whether AI prompts, system instructions, model-selection records, output logs, validation notes, or vendor interactions must be preserved and disclosed as part of the public record.

That gap creates four risks.

First, **outsourced opacity**: a public body may use an external AI service to generate analysis or drafting assistance while arguing that the underlying reasoning is not truly "its" record.

Second, **record fragility**: prompts, outputs, and system logs may not be retained in a systematic way, making meaningful access impossible even if disclosure is legally justified.

Third, **democratic asymmetry**: the public may be governed by outputs whose provenance, assumptions, and instructions remain hidden.

Fourth, **administrative fairness concerns**: the OIPC has already signaled that AI in the public sector raises issues of fairness, due process, and good governance and should be held to high standards.

The problem is especially serious when AI is used in high-impact contexts such as legislative drafting, policy analysis, regulatory design, electoral boundary analysis, enforcement prioritization, benefits administration, or any other function capable of affecting rights, representation, services, or public obligations.

Reform Objectives

The Government of Alberta should amend its legal and policy framework to achieve six objectives. Ensure that AI-assisted government work remains subject to public accountability.

Confirm that prompts, instructions, model identity, outputs, and reliance records are government records where used for public business.

Prevent destruction or non-creation of records necessary to explain how AI shaped public decisions.

Require proactive disclosure for significant public-sector AI uses.

Preserve legitimate exemptions for privacy, security, cabinet confidentiality, and solicitor-client privilege without creating a blanket “AI secrecy” category.

Strengthen public trust through auditable, reviewable, and searchable disclosure practices.

Core Recommendations

1. Amend ATIA to expressly define AI-generated and AI-assisted records as records subject to access

ATIA should be amended to state that, where a public body uses an automated system or artificial intelligence system in carrying out public business, the following are deemed to be records in the custody or under the control of the public body if created, received, relied upon, or retained for that purpose:

- prompts, instructions, system messages, templates, or parameter settings;
- outputs, drafts, summaries, analyses, rankings, recommendations, or generated text;
- records identifying the system, vendor, version, and deployment context;
- records showing review, validation, revision, acceptance, rejection, or reliance by officials;
- records of contractual terms governing retention, inspection, audit, and access.

This amendment would build directly on Alberta’s existing position that records created by outside contracted consultants may be under a public body’s control, particularly where the public body has relied on them substantially or where contract terms permit inspection or copying.

2. Create a statutory duty to preserve AI interaction records

ATIA and related records-management rules should require public bodies to preserve sufficient records to explain material AI use in public business. A public body should not be permitted to avoid access by failing to retain prompts, outputs, or model-use logs.

Minimum retention should include:

- the date and purpose of use;
- the name and version of the system used;
- the prompt or instruction set used for substantive outputs;
- the output received;
- any material edits by officials;
- the decision about whether and how the output was relied upon;
- any fairness, risk, or validation review conducted.

This aligns with Alberta's current guidance that public bodies must create and maintain accurate records documenting decisions, actions, and communications in the request process, and with POPA-related guidance requiring documented policies, oversight, and internal controls for AI-related uses of personal information and derived data.

3. Establish a public AI registry for Alberta public bodies

Alberta should create a publicly accessible and searchable **Public Sector AI Registry**. The OIPC has already recommended that Alberta consider regulations or policy requiring public bodies to be transparent regarding AI use through a searchable public registry.

For each significant AI deployment, the registry should include:

- the public body using the system;
- the vendor or model name;
- the general purpose of use;
- whether the system is external or internally hosted;
- whether personal information or confidential government information is entered into the system;
- the legal authority for the use;
- whether the system supports drafting, analysis, recommendations, classification, decision-making, or communications;
- whether the use is high-impact;
- whether a privacy or fairness assessment was completed;
- whether human review is required before use or publication.

The registry need not expose sensitive technical details that would create security risks, but it should disclose enough to let the public, journalists, auditors, and legislators identify where AI is shaping government action. That is consistent with the OIPC's recommendation and with Alberta's broader privacy-management approach that emphasizes accountability and documentation.

4. Require proactive disclosure for high-impact AI uses

For high-impact uses, public bodies should not wait for an access request. They should proactively disclose a standard transparency package.

High-impact uses should include, at minimum:

- legislative drafting support;
- policy recommendations to Cabinet or ministers, unless shielded by an existing lawful exception;
- electoral or boundary analysis;
- benefits, licensing, inspection, or enforcement prioritization;
- systems that materially affect rights, eligibility, burdens, representation, or public entitlements.
- The proactive disclosure package should include:
 - a plain-language description of the use;
 - the name of the tool or model;
 - the categories of prompts or instructions used;
 - a summary of human review and validation;
 - any privacy, security, or fairness assessment completed;
 - the date range of use;
 - a contact point for requesting further records under ATIA.

This would complement, not replace, ordinary access rights.

5. Add procurement and contract clauses to preserve public access

Any contract for external AI services used by Alberta public bodies should require:

retention of prompt and output records for a defined period;
the public body's right to inspect, copy, export, and disclose records as required by ATIA;
prohibition on contractual clauses that prevent lawful disclosure;
logging of system version changes and material configuration changes;
clear terms on whether government information is used to train vendor systems;
auditing rights where the system is used in high-impact contexts.

This is a practical extension of Alberta's own ATIA guidance, which already identifies contractual inspection and copying rights as indicators that records are under the public body's control.

6. Clarify that existing exemptions apply narrowly and record-by-record

The reform should expressly state that AI-related records are not categorically exempt. Existing exemptions may still apply where justified, including:

personal privacy;
solicitor-client privilege;
security-sensitive information;
cabinet confidentiality where lawfully applicable;
trade secrets or confidential business information, but only to the extent genuinely necessary.

The default should remain disclosure. A vendor's commercial interest should not be allowed to conceal the fact that government relied on a particular system or the substance of prompts and outputs that shaped public action, except where withholding is strictly necessary under existing law.

7. Create a review right focused on AI transparency refusals

Applicants should have an express right to seek expedited review by the OIPC where a public body:

- refuses to confirm whether AI was used;
- claims that prompts or outputs are not records;
- fails to preserve relevant AI records;
- over-redacts AI-related materials using generalized vendor-confidentiality claims.

This is consistent with the OIPC's oversight role under ATIA and with its ongoing work emphasizing the need to modernize Alberta's framework in light of AI and related technologies.

Proposed Legislative Language

The following language is offered as a model drafting direction, not final legislative text.

Proposed new section: Records relating to automated systems and artificial intelligence

For the purposes of this Act, where a public body uses an automated system or artificial intelligence system in relation to a program, service, policy, recommendation, analysis, draft, or decision, any record created, received, used, relied upon, or retained in connection with that use is deemed to be in the custody or under the control of the public body.

Without limiting subsection (1), such records include:

- a. prompts, instructions, system messages, templates, parameter settings, and input materials;
- b. outputs, responses, draft text, summaries, analyses, rankings, classifications, and recommendations;
- c. records identifying the system, model, vendor, version, and purpose of use;
- d. records of human review, validation, revision, approval, rejection, or reliance;
- e. contractual records respecting retention, audit, inspection, copying, confidentiality, and disclosure.

A public body must create and retain records sufficient to document any material use of an automated system or artificial intelligence system in carrying out public business.

No record described in this section is exempt from this Act solely because it was generated, assisted, hosted, processed, or stored by an external service provider or automated system.

Proposed regulation: Public Sector AI Registry

A regulation should require every public body to maintain and publish prescribed information concerning each significant AI use, including purpose, legal authority, system identity, data categories used, assessment status, and oversight mechanisms.

Implementation Plan

Phase 1: Policy and procurement changes within 6 months

The government should immediately issue an interim directive requiring ministries and public bodies to:

- retain prompts and outputs for material AI uses;
- include ATIA-compliant disclosure clauses in all new AI-related contracts;
- identify existing high-impact AI uses;
- designate an internal AI transparency lead.

This can begin before legislative amendment.

Phase 2: Legislative amendments within 12 months

Amend ATIA and related regulations to:

- define AI interaction records;
- establish the preservation duty;
- create registry authority;
- authorize prescribed proactive disclosure categories.

Phase 3: Operational rollout within 18 months

Public bodies should then:

- publish registry entries;
- update records schedules;
- train ATI staff and program staff;
- adopt templates for AI-use documentation;
- integrate AI-use logging into privacy management and information governance practices.

This would fit with Alberta's existing move toward privacy management programs and documented internal AI-related policies for public bodies handling sensitive or high-volume information.

Why This Reform Is Necessary Now

Alberta already has official policy infrastructure acknowledging government AI use. The province's policy tools portal lists an **Artificial Intelligence Usage Policy** described as enabling government staff to use AI in a transparent, responsible, secure, ethical, and human-centered manner, alongside related references to a Data Ethics Framework and Privacy Management Framework.

At the same time, the OIPC has publicly warned that Alberta's new access framework has diminished some access rights through additional exceptions, broader disregard powers, and longer

extensions, and has emphasized the importance of modernizing the framework to deal with technologies such as AI.

Once AI is in use, transparency cannot be left to internal goodwill. It must be guaranteed by law.

Conclusion

Alberta's access-to-information regime should be updated to make one principle unmistakable:

If artificial intelligence is used to help shape public law, public policy, public analysis, or public decisions, then the identity of the system, the prompts that materially directed it, the outputs it produced, and the records showing how officials relied on those outputs must be treated as public records, subject only to narrow and existing exceptions.

That reform would not hinder innovation. It would place innovation inside democratic accountability, where it belongs.

Anticipated Objections and Responses

“This will chill government use of beneficial AI tools.”

The opposite is more likely true. Transparency requirements create institutional discipline around AI use, which reduces legal and reputational risk. Governments that have adopted AI registries and disclosure frameworks — including in the EU and parts of the UK — have not retreated from AI adoption; they've adopted it more deliberately. Where a ministry cannot document why a particular AI tool was used for a high-impact function, that use should be presumed inconsistent with basic standards of accountable public administration. The proposal doesn't prohibit AI use — it requires that use to be legible.

“Prompts and instructions are proprietary vendor information.”

This conflates two different things. A vendor's underlying model architecture, training data, and system design may be genuinely proprietary. But the prompts a *government official* writes to direct that system in carrying out public business are government work product, not vendor trade secrets. The proposal is careful to target records created, received, or relied upon by the public body — not the vendor's internal systems. To the extent vendors contractually claim ownership over government-authored prompts, that is itself a governance problem the procurement clause recommendation addresses directly.

“The compliance burden will be unworkable, especially for smaller agencies.”

The proposal already tiers its obligations — routine uses face lighter requirements than high-impact ones. That said, it could go further by recommending standardized logging templates and shared infrastructure so smaller public bodies aren't building compliance systems from scratch. The cost argument also needs to be weighed honestly against the status quo: the alternative isn't zero cost, it's

the cost of litigation, OIPC investigations, and public trust erosion when AI-assisted decisions are challenged without adequate records.

“Existing exemptions already cover sensitive cases.”

They do, but that’s not the objection it appears to be. The proposal doesn’t remove or override any existing exemption — cabinet confidentiality, solicitor-client privilege, personal privacy, and security-sensitive information all remain intact. The problem the proposal targets is different: the risk that AI-related records will be withheld not under a recognized exemption, but under the threshold argument that they aren’t “records” at all, or aren’t “in the custody or control” of the public body. That gap is real and the existing exemption structure doesn’t close it.

“Public disclosure of prompts could reveal government strategy or expose security vulnerabilities.”

Legitimate concern, already accommodated. Prompts that contain cabinet-confidential deliberations remain protected under that exemption. Prompts that reveal security system logic can be withheld on security grounds. The proposal specifically states that AI records are subject to *existing* narrow exceptions, not a new blanket exemption. What it prevents is using vague security or confidentiality concerns as a catch-all to avoid any disclosure about AI use — including the basic facts of which systems were used and for what purpose.

“The OIPC doesn’t have the capacity to handle a new category of complaints.”

This is the most practically serious objection, and it deserves an honest answer. The proposal does add to the OIPC’s mandate by creating an express review right for AI transparency refusals. The government should accompany these reforms with a resources commitment to the OIPC. It’s worth noting, though, that the OIPC itself has called for these reforms — it is not asking to be shielded from this work. And a well-designed registry and proactive disclosure regime should reduce complaint volume by answering many questions before they become disputes.

About the Author

Lawrence Nault is a Canadian author, filmmaker, and independent researcher based in Alberta's Badlands and Rockies. He writes speculative fiction and YA fantasy, directs documentary film, hosts the podcast *Stone & Signal*, and publishes academic research and policy proposals on digital infrastructure, occupational health, and environmental governance. His work asks hard questions about ecology, technology, and what it means to be human.

Contact

lawrencenault.me/contact/

Website

lawrencenault.me

Policy proposals

lawrencenault.me/policy/

Research

lawrencenault.me/research-policy/

Social

Threads: [@mountainhermitauthor](https://www.threads.net/@mountainhermitauthor)

Bluesky: [@mountainhermit.bsky.social](https://bsky.app/profile/mountainhermit.bsky.social)