



Submission to the Public Consultation on Advancing AI Transparency in Canada

Respondent: Raval Consulting Inc., Richmond Hill, Ontario

Submitted by: Kapil Raval, Managing Principal

Date: August 2, 2026

Re: *Enhancing trust in artificial intelligence through increased transparency* — Discussion Paper

1. About This Submission

Raval Consulting Inc. is a technology advisory firm based in Richmond Hill, Ontario, with a focus on AI adoption, commercialization, governance, and partner ecosystem strategy. The firm's work spans three key stakeholder groups within the AI ecosystem: technology vendors developing AI-enabled products and services, channel and implementation partners responsible for deployment, and organizations evaluating the adoption of AI solutions.

My own work sits in a specific part of that ecosystem. I advise the vendors selling AI-enabled solutions into the Canadian mid-market, the partner networks through which those sales actually happen, and the organizations on the other side of the table deciding whether to adopt. Current engagements span Canada, Europe, and Asia, which gives me a working view of how transparency requirements in other jurisdictions land on the companies that have to comply with them.

I established the practice in 2018, initially advising the boards of two Canadian AI companies that had recently completed public listings, and served in that capacity for three years. I set it aside as my responsibilities at Microsoft grew and returned to it full time after retiring from corporate life. Before that I spent more than twenty-five years in senior global roles at Microsoft, Hewlett Packard Enterprise, Compaq, Digital Equipment Corporation, and Philips, including regional leadership across the Americas, Asia-Pacific, and the Middle East. Over that period I watched the enterprise software industry attempt to build trust through transparency, compliance, and partner programs several times over. Some of those attempts worked. Others did not, and they failed in instructive ways.

This submission does not represent the perspective of a model developer or platform operator, though board work gave me a close view of what disclosure and compliance obligations cost a Canadian AI company to meet. It reflects the perspective of the organizations that evaluate, purchase, deploy, and govern AI-enabled systems. Transparency policy succeeds or fails at the point of procurement and deployment rather

than during development, and that is the vantage point I am writing from: sitting across a table from a mid-sized Canadian company as it decides whether to buy something.

Accordingly, this submission focuses on practical measures that can strengthen trust, accountability, and adoption outcomes while minimizing unnecessary compliance burdens for Canadian businesses.

2. Summary of Recommendations

Guiding Principle

The recommendations in this submission are guided by a single principle:

Transparency obligations should be calibrated to the consequences and risk profile of a system's use rather than the mere presence of artificial intelligence.

The distinction is important, and it matters commercially. A synthetic video of a public figure saying something they never said is a problem because it is false, not because a model produced it. A proposal drafted with AI assistance and reviewed by a competent professional is not a problem at all. If the trigger for disclosure is AI involvement, Canadian businesses will be required to label ordinary work product at real cost, with no safety benefit, and with the unintended message that using AI is something to confess. That works directly against the adoption the strategy is trying to encourage.

A risk-based approach also helps distinguish between **provenance** and **warning**. Provenance mechanisms provide useful information about how content or decisions were generated. Warning mechanisms should be reserved for situations involving a material risk of deception, misrepresentation, or confusion regarding whether a person is interacting with a human or an automated system.

R1. Strengthen Buyer Readiness

The primary barrier to responsible AI adoption among many Canadian mid-market organizations is not a lack of information from developers but a lack of organizational readiness. Transparency initiatives should place greater emphasis on enabling organizations to assess their own preparedness to deploy AI safely and effectively.

R2. Develop the SME Readiness Assessment Tool in Collaboration with End Users

The AI for All strategy's proposed AI Literacy and Adoption Assessment tool represents a valuable opportunity to strengthen adoption outcomes. To be effective, it should be designed for business decision-makers rather than technical specialists, require minimal time to complete, and address organizational readiness alongside vendor evaluation.

The tool should be tested extensively with SME operators, advisors, channel partners, and business leaders prior to release.

R3. Explicitly Recognize Intermediaries in the AI Value Chain

Most Canadian organizations do not obtain AI solutions directly from model developers. They acquire them through resellers, systems integrators, managed service providers, software vendors, and consultants.

Any transparency framework that focuses exclusively on developers and end users risks overlooking the point at which most procurement and deployment decisions occur.

R4. Establish Transparency Obligations Throughout the Distribution Chain

Transparency requirements should flow contractually through the AI value chain. Organizations responsible for implementing, configuring, or reselling AI systems should be required to communicate relevant disclosures and implementation-specific modifications to customers.

R5. Base Obligations on Use and Consequence Rather Than Organizational Size

A 20-person firm putting AI in control of warehouse robotics can injure someone. A 2,000-person firm using AI to summarize its own internal meetings can produce a poor summary. Reverse the employee counts and the risks stay exactly where they were, because they belong to the application rather than to the organization. Size is a reasonable proxy for a company's capacity to comply; it is a poor proxy for the consequences of failure, and the two should not be conflated. Canada already regulates this way in other domains — a small machine shop is not exempt from workplace safety rules because it is small.

R6. Establish Foundational Requirements for AI Agents

Government should move quickly to establish baseline expectations regarding AI agents while avoiding highly prescriptive technical requirements. A workable definition is available now: a system that takes an action without a person approving each step. The distinction that matters is not how autonomous an agent is but whether its actions can be undone — an agent that drafts a document and one that moves inventory, dispatches a vehicle, or releases a payment carry very different consequences.

At a minimum, AI agents should be clearly defined, their actions should be auditable, and records should indicate what actions were taken, when, under whose authority, and within what operational constraints.

R7. Prioritize Measurable and Verifiable Obligations

Transparency requirements should be drafted in ways that permit objective verification.

“Deployers shall be transparent about their use of AI” cannot be verified by anyone. “On request, name the person accountable for this system and state whether its outputs can be traced to a source” can be verified in a phone call. Sarbanes-Oxley is the cautionary case: it certified that controls existed, not that they worked, and became a paperwork exercise as a result.

R8. Use Government Procurement as a Policy Lever

Federal procurement policies can influence market behavior more rapidly and effectively than legislation alone. Transparency requirements incorporated into public-sector procurement processes can establish commercial expectations that extend well beyond government contracts.

R9. Align Incident Reporting with Severity of Harm

Incident-reporting obligations should be triggered by the seriousness of potential or actual harm rather than by AI involvement alone. A consequence-based approach will improve reporting quality while reducing unnecessary burden on smaller firms.

R10. Provide Practical Guidance for Organizations Handling Public-Interest Information

Community organizations, trade associations, advocacy groups, and smaller businesses increasingly use AI to summarize policy, develop public-facing content, and interpret information. Most operate without dedicated governance expertise.

Plain-language guidance, supported by safe-harbour provisions for organizations that follow recognized good practices, would likely deliver greater public benefit than additional compliance obligations.

R11. Require Systems to Demonstrate the Basis for Their Outputs

Transparency is most valuable at the point of decision-making. AI systems should show users the source material an output rests on, so that a person can catch an error before acting on it. Unlike documentation, this can be tested, which makes it suitable both for government procurement and for any buyer’s evaluation checklist. Where a system cannot show its basis, that limitation should be stated plainly at the outset rather than left for the user to discover.

R12. Require Transparency Regarding Training Data Sources and Liability

Organizations purchasing AI systems have a legitimate interest in understanding:

- The categories and provenance of training data used to develop a system;
- The intellectual property and licensing conditions associated with that data; and
- Which party bears responsibility for resulting legal or compliance risks.

Standardized disclosures regarding training data and indemnification arrangements would improve market transparency and support informed purchasing decisions.

R13. Design the Canada Trusted AI Certification Program Against Known Failure Modes

The certification program announced in AI for All is the most consequential transparency measure currently in development, and its scope, evidentiary basis, and enforcement mechanism have not yet been defined. Four requirements follow from the record set out in Section 3. Certification should rest on independent verification rather than supplier self-attestation. It should carry a working complaint and remediation path, because an audit obligation that nobody can invoke produces no accountability. It should be structured for periodic re-validation rather than one-time approval, because a certified model can be retrained the week after it is certified. Finally, the cost and duration of certification should be scaled so that it does not become a moat protecting incumbents. FedRAMP is the cautionary precedent: a substantive assurance regime whose expense and timelines effectively excluded smaller vendors from United States federal cloud procurement, raising the barrier to entry rather than the standard of the market. A certification tier that a 20-person Canadian firm cannot afford to enter would work against the domestic scale-up objectives the strategy sets out elsewhere.

Certification schemes are difficult to repair once buyers have learned to treat them as a formality. The design stage is the point at which this is inexpensive to get right.

3. Lessons from Earlier Transparency Regimes

Transparency and assurance requirements are not new. Enterprise technology markets have experimented with a wide range of trust-building mechanisms over several decades, yielding lessons directly relevant to AI governance.

Trust Is Often Mediated Through Intermediaries

Through the 1980s and 1990s, enterprise software was sold through value-added resellers, and what the buyer trusted was not the vendor's documentation but the local firm that installed the system and answered the phone when it broke. That trust was real, and it had two defects: nobody outside the room could verify it, and it did not scale.

The industry spent three decades trying to replace that relationship with certificates and audit reports, and did not succeed. Even now, the intermediary is where the real commercial relationship lives, regardless of how much formal paperwork the large vendors produce. A framework aimed only at the original developer and the final user targets two parties who almost never speak to each other.

A practical solution is a pass-through obligation requiring implementation partners to provide customers with developer disclosures as well as a clear description of any material configuration or customization performed during deployment.

Self-Attestation Alone Does Not Produce Trust

The EU-US Safe Harbor framework, its successor Privacy Shield, and the TRUSTe seal programme all relied on organizations certifying their own compliance. All three failed: two were struck down by the courts, and the third drew regulatory action for not carrying out the reviews it advertised. Voluntary AI commitments are built the same way.

Verification is no longer hypothetical. The European Union requires conformity assessment for high-risk systems, ISO 42001 offers accredited certification, and Canada already applies an Algorithmic Impact Assessment to its own automated decision systems.

Verification on its own is not sufficient either. New York City's Local Law 144 requires an independent bias audit of automated hiring tools and publication of the results; a study of 391 employers found fewer than twenty had posted one, and the city's comptroller later found that most complaints never reached the enforcing agency. An audit obligation works only where someone can raise a complaint and someone is required to act on it.

Disclosure Without Capability Does Not Deliver Transparency

Information has value only when recipients can understand and act upon it.

Cookie consent banners increased the volume of disclosure while reducing comprehension. Software bills of materials, mandated after recent supply chain incidents, are now produced at scale and consumed by very few. AI transparency measures should be designed backwards from the party expected to use the information, and tested on that party before release.

This consideration is particularly important for SMEs and non-technical decision-makers.

Assurance Must Keep Pace with Technological Change

Security certification schemes such as the Orange Book and, later, Common Criteria required evaluation cycles measured in years against products that shipped in months. Buyers stopped treating the certificates as meaningful long before the schemes were formally retired.

Given the pace of AI development, transparency and assurance mechanisms should emphasize periodic re-validation and continuous monitoring rather than relying exclusively on point-in-time certification. The industry has already made this shift once: zero-trust security replaced the assumption that a system checked at the perimeter stays trustworthy with continuous verification of every request, and it is now mainstream enterprise practice. The same reasoning applies to a model that can be retrained or updated between assessments.

Key Implication

Across multiple generations of technology regulation, transparency measures have proven most durable when they reduce decision-making costs for both the disclosing party and the receiving party. Conversely, transparency initiatives tend to become compliance exercises when they impose costs without creating practical utility.

This principle supports the emphasis throughout this submission on measurable obligations, buyer readiness, and decision-time transparency. It also applies directly to a measure already announced. The Canada Trusted AI Certification program is a certification scheme in the early stages of design, and the four findings above describe, with some precision, the ways in which comparable schemes have previously failed (R13).

4. The Reality for Small and Medium-Sized Enterprises

The discussion paper appropriately recognizes the importance of SMEs within the Canadian economy. However, many of the proposed measures focus primarily on increasing the flow of information from developers to organizations and from organizations to the public.

This approach assumes that recipients possess the expertise needed to interpret and act upon the information they receive.

In many Canadian mid-market organizations, that assumption does not hold.

The Statistics Canada figures cited in the discussion paper show that roughly one in five Canadian businesses used AI in the second quarter of 2026, up from about one in eight a year earlier and triple the proportion reported in 2024. This growth is not occurring in companies with a Chief AI Officer or a data science team. It is occurring in organizations of 50 to 500 people, where the purchasing decision is made by a general manager, often on the strength of a vendor's sales presentation.

In my advisory work, the questions that determine whether a deployment succeeds are consistently organizational rather than technical:

- Is there a clearly defined business case?
- Are data assets fit for purpose?
- Is accountability clearly assigned?
- Is there capacity to monitor performance after deployment?
- Are governance processes adequate for the level of risk involved?

Organizations unable to answer these questions are unlikely to benefit from highly technical system documentation, regardless of its quality.

For this reason, strengthening buyer capability should be considered a core transparency objective. The proposed AI Literacy and Adoption Assessment tool offers a practical

mechanism for achieving this goal if developed with meaningful input from those responsible for implementing AI within SMEs.

5. Information About AI Systems (Section 3.3)

The discussion paper correctly notes that existing AI documentation is often highly technical and inaccessible to non-specialist audiences.

The challenge extends beyond readability. Much of the current documentation ecosystem was designed for technical practitioners rather than business decision-makers.

A useful transparency framework should recognize three distinct audiences:

Technical Teams

Technical professionals require detailed information regarding model performance, testing methodologies, security controls, and limitations.

Business Decision-Makers

Business leaders require information that supports procurement and governance decisions, including:

- Appropriate use cases;
- Material limitations;
- Key risks;
- Human oversight requirements; and
- Consequences of system failure.

Individuals Affected by AI Decisions

Individuals interacting with or affected by AI systems need to know:

- Whether AI played a role in a decision;
- What recourse mechanisms exist; and
- How decisions can be challenged or reviewed.

At present, responsibility for serving the needs of business decision-makers is poorly defined.

Transparency is particularly valuable at the point of decision-making. The discussion paper's example of fabricated legal citations shows why. The failure there was not that AI involvement went undisclosed; disclosure would have changed nothing. It was that the system asserted something with no traceable basis, leaving the user no way to notice before acting. No quantity of upfront documentation prevents that. Showing the source would have.

The same principle applies to transparency regarding training data and liability. Organizations purchasing AI systems should understand both the origin of training data and the allocation of legal responsibility associated with its use.

Similar challenges exist among smaller organizations producing public-facing content using AI tools. Most lack access to specialized governance expertise and would benefit substantially from practical guidance regarding responsible use.

6. The Missing Layer: Intermediaries

The discussion paper focuses primarily on developers, deployers, and users.

An important stakeholder group receives comparatively little attention: the intermediaries responsible for bringing AI solutions to market and supporting their implementation.

For many Canadian organizations, the most significant commercial relationship is not with the model developer but with a reseller, software provider, managed service provider, systems integrator, or consultant.

This creates two important challenges.

First, information may not flow consistently from developers to end customers.

Second, responsibility can become fragmented when multiple organizations contribute to the development, implementation, customization, and operation of a system.

Effective transparency frameworks should therefore explicitly recognize intermediaries and require disclosure obligations to flow throughout the supply chain.

The goal should be to ensure that customers receive relevant information from the organization with which they maintain a direct commercial relationship.

7. AI Agents (Section 3.5)

AI agents represent one of the most rapidly evolving areas of AI deployment.

Government should establish foundational governance principles while avoiding highly prescriptive technical requirements that may become obsolete quickly.

A useful distinction can be drawn between:

- Systems that generate outputs for human review; and
- Systems that autonomously initiate actions in operational environments.

The second category carries materially greater risk. Once an agent can operate equipment, adjust a physical process, or commit funds, a failure is not a poor recommendation but an event that has already happened, and the transparency has to be available before the action rather than after it. Canadian manufacturing, logistics, and field service are where this shift is now underway, often in firms too small to have anyone assigned to think about it.

Two foundational principles warrant early consideration.

Identification

Individuals interacting with an AI agent should be able to determine clearly that they are engaging with an automated system and identify the organization responsible for its operation.

Record-Keeping

Where agents perform significant actions, robust audit records should capture:

- What action occurred;
- When it occurred;
- What information was accessed;
- Under whose authority the action was taken; and
- What operational permissions were in effect.

Organizations should also receive clear, accessible descriptions of an agent's capabilities, permissions, and constraints at the point of deployment.

8. Using Government Procurement to Shape Market Behaviour

Government procurement represents one of the most effective instruments available to influence industry practices.

History demonstrates that procurement standards frequently establish broader commercial norms in areas such as accessibility, cybersecurity, and risk management. Accessibility is the clearest case: United States federal procurement rules under Section 508 made accessible design a commercial default worldwide, because no vendor was going to maintain one accessible product for government buyers and a separate one for everyone else.

Canada already applies impact-based governance principles through the Treasury Board Directive on Automated Decision-Making and its associated Algorithmic Impact Assessment framework.

Extending comparable transparency expectations to government suppliers would build upon existing practices rather than creating an entirely new regulatory model.

Procurement-led approaches also provide a practical mechanism for encouraging adoption of good practices without imposing unnecessary burdens on organizations that do not participate in public-sector markets.

9. Incident Reporting (Section 3.4)

The discussion paper's exploration of aviation-style incident reporting is constructive and merits further consideration.

The principal policy question concerns threshold design.

Reporting requirements should focus on events involving meaningful harm, material risk, or significant public-interest consequences.

Consider one company running two systems. Its AI scheduling assistant double-books meetings across a client base for a week: a commercial annoyance, and nobody learns anything from a filing about it. Its AI triage tool quietly deprioritizes a category of patients: same company, same headcount, same underlying technology, and a genuine reportable event. Aviation already draws the line this way, since a near-miss report is triggered by proximity to harm rather than by which system was involved.

A consequence-based approach will improve signal quality, reduce unnecessary administrative burden, and encourage organizations to focus attention on events from which meaningful lessons can be learned.

10. Conclusion

Canada has an opportunity to distinguish itself not by being the first jurisdiction to introduce AI transparency requirements, but by implementing measures that are practical, proportionate, and effective.

Much of the current policy discussion appropriately focuses on the responsibilities of AI developers and model providers. However, the long-term success of any transparency framework will ultimately depend on whether organizations deploying AI systems are able to make informed decisions regarding procurement, implementation, governance, and oversight.

Particular attention should be paid to the realities of the Canadian mid-market, where adoption decisions are frequently influenced by channel partners, systems integrators, managed service providers, and software vendors rather than direct relationships with AI developers.

A successful framework should prioritize:

1. **Transparency obligations proportionate to the consequences of system use;**
2. **Practical measures that strengthen buyer readiness and organizational capability;**
3. **Clear accountability throughout the entire AI value chain, including intermediaries; and**
4. **Transparency mechanisms that support decision-making rather than paperwork generation.**

If transparency measures enable organizations to adopt AI responsibly, understand their obligations, evaluate risk effectively, and hold suppliers accountable, Canada will advance both public trust and economic competitiveness while avoiding unnecessary compliance burdens.

Thank you for the opportunity to contribute to this important consultation.

Raval Consulting Inc.

Richmond Hill, Ontario

kapil@ravalconsulting.com | ravalconsulting.com