

Canada's new AI for All strategy: A business outlook on AI governance, adoption, and data sovereignty

June 09, 2026

Canada's release of its national artificial intelligence strategy, AI for All, marks a substantive shift in federal policy, signalling how the government intends to govern AI for the foreseeable future.

The strategy's most significant feature is what it is not: as previously announced, it does not revive the *Artificial Intelligence and Data Act* (AIDA), the proposed omnibus AI statute that stalled in Parliament and was effectively abandoned following the change in government earlier this year.

Rather than returning to a centralised, risk-based legislative regime, the federal government has chosen a different path: a distributed governance model that combines targeted legal reform, public investment, sovereign infrastructure, and continued reliance on existing frameworks — privacy, consumer protection, human rights, and sectoral regulation — as the primary tools for managing AI risk.

For Canadian businesses, that choice has immediate practical consequences. AI governance is not on pause while Ottawa designs a new statute. Existing legal obligations apply today, across privacy law, human rights legislation, consumer protection, and sector-specific regulation. What the strategy adds is a clearer articulation of where those frameworks are heading, and a set of economic and industrial priorities that will shape how regulators, funders, and procurers engage with AI over the next several years.

What's new: Four meaningful objectives from AI for All

AI for All is organised around six pillars, but for Canadian businesses, four operational elements carry the most immediate significance.

Adoption as a primary policy objective

The strategy sets explicit national targets: \$200 billion in additional economic growth, 250,000 new AI-related jobs over five years, and an increase in AI adoption from roughly 12 per cent to 60 per cent by 2034.

These figures are aspirational rather than legally binding, and they should be read as such. What matters for businesses is not the numbers themselves, but what they signal: the federal government has made adoption at scale a central policy priority, and funding decisions, procurement criteria, and regulatory posture will increasingly reflect that priority.

The strategy also positions government as an active adopter of AI systems, not merely a regulator. Public-sector deployment and procurement are explicitly framed as mechanisms for setting expectations around trusted and compliant AI, meaning that companies seeking government contracts will likely face governance requirements that precede any formal legislative mandate.

Sector-driven deployment

The strategy identifies priority sectors for accelerated AI adoption: health and life sciences, energy and natural resources, transportation, agriculture, and manufacturing. For businesses operating in these areas, the practical implication is that AI deployment is increasingly likely to occur through structured federal initiatives, targeted funding programs, and public–private partnerships, not purely through internal innovation cycles.

The strategy also signals a more active role for government as an adopter and procurer of AI systems, using public-sector deployment to establish expectations for trusted and compliant AI. This matters legally because participation in government-supported programs typically carries conditions: around data governance, interoperability, procurement rules, and accountability.

Organizations in priority sectors should assess those conditions carefully before assuming that federal support for adoption is unconditional. The health sector provides an early example, with initiatives such as VITAL using federated data models that embed specific governance requirements directly into the architecture of AI deployment.

Sovereign infrastructure as policy

One of the strategy's most consequential elements is its treatment of AI infrastructure as a matter of national policy. Building on the Canadian Sovereign AI Compute Strategy, the federal government is investing in domestic supercomputing capacity, data-centre infrastructure, and expanded access to compute resources, signalling that it expects Canadian organizations to take seriously where their AI systems are built, trained, and hosted.

For businesses, this means that decisions about cloud providers, data residency, and cross-border data flows are no longer purely technical or commercial choices. They are increasingly strategic and, in some contexts, regulatory ones. Organizations whose AI systems rely heavily on foreign-controlled infrastructure or data arrangements that are difficult to reconcile with Canadian governance expectations should treat this as a

material risk management issue, particularly in regulated sectors or where government procurement is relevant.

Trust through distributed governance

The strategy's approach to AI risk and accountability deserves close attention, precisely because it does not take the form of a single statute. Instead, the federal government is signalling increased intervention across multiple existing legal channels: privacy law modernisation, online safety regulation, measures targeting deepfakes and surveillance pricing, and an expanded mandate for the Canadian AI Safety Institute.

Of particular note is the proposed Canada Trusted AI Certification program, intended to identify trustworthy AI products in the marketplace. The legal significance of this mechanism will depend entirely on its design: whether certification is voluntary or effectively mandatory, whether it creates safe harbours from regulatory scrutiny, and whether it becomes a condition of public procurement or sector-specific licensing. Businesses should monitor this closely: certification regimes that begin as voluntary frequently become baseline expectations in regulated procurement and high-stakes deployment contexts.

The practical effect is a distributed compliance model in which privacy, consumer protection, cybersecurity, and sector oversight increasingly converge on AI systems. The absence of a single AI statute does not reduce compliance complexity; in many respects, it increases it, because organizations must track and reconcile obligations across multiple frameworks simultaneously.

Taken together, these four elements reflect a coherent, if demanding, policy direction: accelerate adoption, anchor infrastructure and data domestically, and embed accountability expectations across multiple legal and regulatory channels. For businesses, the result is an environment that is more interventionist than the absence of omnibus legislation might suggest. The compliance burden has not been deferred: it has been distributed.

Global context: A fragmented regulatory landscape on artificial intelligence

Canada's AI for All strategy must be understood within a global environment that is not merely fragmented, but actively diverging, with major jurisdictions now pursuing fundamentally different philosophies on AI governance.

AI in the EU

The European Union remains the most consequential foreign regime for Canadian businesses. The EU AI Act is a comprehensive, risk-based framework with explicit extraterritorial reach: Canadian organizations that develop or deploy AI systems whose outputs affect individuals in the EU may be in scope regardless of where they are incorporated, or where their systems are hosted.

For many Canadian businesses, EU compliance is not a future consideration, it is a current legal obligation that requires attention now, including conformity assessments, transparency requirements, and in some cases prohibited-use restrictions that apply irrespective of Canadian law.

However, by adopting a sector-led model, relying on existing regulators rather than a single statute, the United Kingdom currently sits closest to Canada's chosen approach.

AI in the U.S.

The United States presents a sharply contrasting picture, and one that shifted materially on June 2, 2026, as President Trump signed an executive order titled "Promoting Advanced Artificial Intelligence Innovation and Security."

The directive orders federal agencies to establish a framework for the secure deployment of frontier AI models, including a voluntary process by which developers would provide the government with early access to models for up to 30 days before broader release. The order attempts to shore up the country's cyber defences without compelling AI companies to share information about their latest systems. An earlier draft required a 90-day government review window before model release; that timeline was cut to 30 days in the final order, following significant industry lobbying over concerns about competitive harm.

The practical implication for Canadian businesses is significant. The U.S. is now explicitly pursuing a deregulatory, innovation-first posture on AI, with voluntary rather than mandatory oversight mechanisms. Canadian organizations competing with or operating alongside U.S. firms will face a structural asymmetry: a more permissive environment south of the border, a more prescriptive one in the EU, and an evolving distributed framework at home. That asymmetry creates both competitive pressure, as U.S. firms may move faster with less governance overhead, and compliance complexity for organizations operating across all three jurisdictions.

Canada's AI for All in the global context

For Canadian businesses, the immediate strategic consequence of this fragmentation is clear: compliance cannot be designed around domestic requirements alone.

- Organizations with any EU market exposure should treat EU AI Act obligations as the compliance floor, not a future consideration.
- Those operating in the U.S. market should monitor how the June 2 executive order develops in practice, particularly whether the voluntary pre-deployment review process becomes an informal condition of federal procurement or partnership.
- All organizations should assume that the gap between jurisdictions will create ongoing pressure to maintain governance frameworks that are scalable and adaptable, rather than jurisdiction-specific.

Canada's position in this landscape — more interventionist than the U.S., less prescriptive than the EU — may prove to be a competitive advantage if the distributed governance model is implemented coherently. The risk is that it creates a compliance gap: not regulated enough to provide the clarity that sophisticated governance

frameworks require, but not deregulated enough to match the speed at which U.S. competitors can deploy.

How Canadian organizations can adjust to AI for All: Key lessons

The strategy's practical impact will be defined by execution, but several legal and regulatory implications are already apparent, and warrant immediate attention.

Governance expectations are increasing, even without an AI statute

The absence of a comprehensive AI statute does not mean the legal risk landscape is undeveloped. Canadian businesses deploying AI systems are already operating within a framework of enforceable obligations, and the strategy signals that enforcement attention in these areas will increase. The most immediate exposure sits in three areas.

First, privacy law: the use of personal data to train, operate, or improve AI systems engages obligations under PIPEDA and provincial equivalents, including requirements around consent, purpose limitation, and the ability to explain automated decisions. The Office of the Privacy Commissioner has already signalled that AI deployments are a priority enforcement area, and the strategy's commitment to privacy modernisation suggests those obligations will become more demanding, not less. Against that backdrop, investments in data governance, documentation, transparency, and explainability are forward-compatible.

Efforts to map data flows, formalize purposes, strengthen consent frameworks, and implement explainability processes for AI systems will position organizations to meet the likely contours of a modernized regime, which is expected to feature:

- stronger individual rights and control over personal information;
- expanded transparency and explainability expectations;
- materially increased enforcement powers and penalties;
- closer alignment with international standards.

Organizations that act now are not getting ahead of the law; they are aligning with where it is already headed.

Second, human rights legislation: AI systems used in hiring, lending, insurance, or service delivery that produce discriminatory outcomes are already vulnerable to challenge under federal and provincial human rights frameworks, regardless of whether the discrimination was intended or understood by the deploying organization.

Third, consumer protection: AI-driven pricing, personalisation, and customer-facing automation are increasingly attracting scrutiny under existing consumer protection frameworks, a risk the strategy explicitly acknowledges through its reference to surveillance pricing measures.

What you can do

Organizations should audit their current AI deployments against these three frameworks now, rather than waiting for AI-specific legislation to define the compliance perimeter.

Government procurement and funding conditions are where adoption expectations acquire legal force

AI for All positions the federal government as an active AI adopter, and public procurement is explicitly framed as a mechanism for setting governance expectations.

For businesses that supply AI systems to government, or that participate in federally supported programs in priority sectors, this is not an abstract policy signal. Procurement criteria, funding program conditions, and partnership agreements will increasingly embed specific requirements around transparency, accountability, data governance, and auditability.

What you can do

Organizations pursuing federal contracts or sector-specific funding should treat governance readiness as a procurement requirement, not a post-award consideration. The proposed Canada Trusted AI Certification program, once operational, may function as a de facto threshold condition in this context.

Infrastructure and data decisions are becoming regulated design choices

The strategy's emphasis on sovereign compute and domestic infrastructure has direct legal implications for how organizations structure their AI systems. Decisions about cloud providers, data residency, and cross-border data flows engage an expanding set of legal and regulatory considerations. These include privacy law requirements around cross-border transfers, potential procurement conditions around data sovereignty, and sector-specific obligations in regulated industries.

What you can do

In practical terms, organizations should be asking a specific set of questions: Where is training data sourced and stored? Where are models trained and hosted? What contractual and jurisdictional protections govern access to that data and those systems? For organizations in regulated sectors, or those seeking government procurement or partnership, these questions are likely to become conditions of eligibility, not merely good governance practice.

Sector-specific legal conditions will emerge rapidly and vary significantly

AI deployment in the strategy's priority sectors — health and life sciences, energy, transportation, agriculture, and manufacturing — will increasingly occur through structured federal programs and public-private partnerships that carry their own legal conditions. These are not uniform.

What you can do

Organizations should expect sector-specific requirements to vary materially and be prepared to include some combination of the following:

- In health and life sciences, federated data governance requirements, patient consent frameworks, and regulatory oversight from Health Canada for AI systems that meet the definition of a medical device under the *Food and Drugs Act*.
- In energy and natural resources, data sharing obligations, interoperability standards, and environmental and Indigenous consultation requirements that attach to infrastructure-adjacent AI deployments.
- In financial services (not a named priority sector, but one where AI deployment is already advanced), the Office of the Superintendent of Financial Institutions (OSFI)'s guidance on model risk management represents the most developed sector-specific risk environment in Canada today, involving the application of existing consumer protection and anti-discrimination frameworks to algorithmic decision-making.

Organizations in these sectors should map the specific legal conditions that will govern their participation in government-supported programs before committing to deployment architectures that may be difficult or costly to adjust.

Cross-border compliance requires a jurisdiction-aware governance framework

As set out in the above section on global context, Canadian organizations face a three-way compliance environment: EU AI Act obligations that may already apply; a deregulatory U.S. posture that creates competitive asymmetry; and an evolving domestic framework. The practical consequence is that governance frameworks designed solely around Canadian requirements will be insufficient for most organizations with international exposure.

The specific implication of the June 2 U.S. executive order is worth noting for Canadian businesses with U.S. market presence or U.S.-based AI supply chains. The voluntary pre-deployment review framework established by the order may evolve into an informal condition of U.S. federal procurement, or national security-sensitive partnerships.

What you can do

Canadian organizations in defence, critical infrastructure, or government-adjacent markets should monitor that development closely, as it could affect the terms on which Canadian AI systems or AI-enabled products are accepted in the U.S. market.

Implementation risk is material and should be reflected in planning horizons

The gap between policy intent and operational reality deserves serious attention. Compute buildout faces energy availability and permitting constraints. Privacy modernisation legislation has not yet been tabled. The Canada Trusted AI Certification program is proposed, not operational. Sectoral initiatives are at varying stages of development.

What you can do

For businesses, this means that the regulatory framework will continue to evolve in ways that are difficult to predict with precision. The appropriate response is not to wait for certainty before investing in governance, but to build governance frameworks that are adaptable: designed to meet current obligations while remaining scalable as requirements develop. Organizations that invest in foundational governance infrastructure now will be materially better positioned to absorb regulatory change than those that treat compliance as a future exercise.

BLG can assist

Canada is not waiting for a single statute to define the rules. Neither should Canadian businesses.

If you would like to discuss how the Canada's new AI for All strategy may affect your organization, or to assess your current AI governance framework in light of these developments, BLG's AI lawyers would be pleased to assist; please reach out to the authors or key contacts below.

By

[Hélène Deschamps Marquis](#), [George R. Wray](#), [Frédéric Wilson](#), [Robert Stefanelli](#)

Expertise

[Information Technology](#), [Cybersecurity](#), [Privacy & Data Protection](#), [Compliance with Privacy & Data Protection](#), [Government & Public Sector](#), [Artificial Intelligence \(AI\)](#), [Technology](#)

BLG | Canada's Law Firm

As the largest, truly full-service Canadian law firm, Borden Ladner Gervais LLP (BLG) delivers practical legal advice for domestic and international clients across more practices and industries than any Canadian firm. With over 800 lawyers, intellectual property agents and other professionals, BLG serves the legal needs of businesses and institutions across Canada and beyond – from M&A and capital markets, to disputes, financing, and trademark & patent registration.

blg.com

BLG Offices

Calgary

Centennial Place, East Tower
520 3rd Avenue S.W.
Calgary, AB, Canada
T2P 0R3

T 403.232.9500
F 403.266.1395

Ottawa

World Exchange Plaza
100 Queen Street
Ottawa, ON, Canada
K1P 1J9

T 613.237.5160
F 613.230.8842

Vancouver

1200 Waterfront Centre
200 Burrard Street
Vancouver, BC, Canada
V7X 1T2

T 604.687.5744
F 604.687.1415

Montréal

1000 De La Gauchetière Street West
Suite 900
Montréal, QC, Canada
H3B 5H4

T 514.954.2555
F 514.879.9015

Toronto

Bay Adelaide Centre, East Tower
22 Adelaide Street West
Toronto, ON, Canada
M5H 4E3

T 416.367.6000
F 416.367.6749

The information contained herein is of a general nature and is not intended to constitute legal advice, a complete statement of the law, or an opinion on any subject. No one should act upon it or refrain from acting without a thorough examination of the law after the facts of a specific situation are considered. You are urged to consult your legal adviser in cases of specific questions or concerns. BLG does not warrant or guarantee the accuracy, currency or completeness of this publication. No part of this publication may be reproduced without prior written permission of Borden Ladner Gervais LLP. If this publication was sent to you by BLG and you do not wish to receive further publications from BLG, you may ask to remove your contact information from our mailing lists by emailing unsubscribe@blg.com or manage your subscription preferences at blg.com/MyPreferences. If you feel you have received this message in error please contact communications@blg.com. BLG's privacy policy for publications may be found at blg.com/en/privacy.

© 2026 Borden Ladner Gervais LLP. Borden Ladner Gervais LLP is an Ontario Limited Liability Partnership.