POLICY BRIEFING NOTE

Report: *Sovereign by Design: Strategic Options for Canadian AI Sovereignty*

March 2026

Authors: Sean Mullin and Jaxson Khan

munk school OF GLOBAL AFFAIRS & PUBLIC POLICY
UNIVERSITY OF TORONTO

Purpose

This policy briefing note is a companion document to *Sovereign by Design: Strategic Options for Canadian AI Sovereignty* (www.aicompetitiveness.ca). It provides a condensed summary of the report's analytical frameworks, vulnerability assessment, and policy recommendations for policy practitioners. For the complete analysis, evidence base, and sourcing, readers should consult the full report. Section references below (e.g., "Section V", "Section VIII.B") refer to the full report.

Part 1: Executive Summary

Artificial intelligence is advancing at an extraordinary speed, reshaping economies, national security, and geopolitical competition. Global AI investment is reaching historically unprecedented levels at hundreds of billions of dollars annually, and the United States now explicitly frames AI dominance as a core national security interest. Canada is the birthplace of modern AI, yet controls neither the major firms that dominate its deployment nor the critical infrastructure that powers it. This dependency is a dangerous vector for coercion, and recent tariff threats, territorial provocations, and "51st state" rhetoric have clarified that the rules-based order Canada once relied upon can no longer be assumed.

AI SOVEREIGNTY: This paper argues that sovereignty in the AI era means freedom from coercion, not digital isolationism or technological self-sufficiency. No country can achieve complete independence across the AI technology stack; the question is how to structure dependencies to preserve choice, reduce foreign leverage, and ensure that Canadian data and infrastructure remain governed by Canadian laws and values. We assess sovereignty across five dimensions (jurisdictional, operational, technological, societal, and economic) and find that AI amplifies threats across all five simultaneously (Section V).

GLOBAL CONTEXT: The United States pursues explicit technological dominance, with American platforms serving as vectors for American jurisdiction and power. China is building a self-sufficient AI stack while exporting open-weight models that embed its own political assumptions globally. The European Union combines regulation with industrial capacity-building. Middle powers face a choice between dependency on foreign AI systems or technological weakness; but coalition-building and hybrid strategies offer a path beyond this binary. Canada must navigate these dynamics, including the impending July 2026 CUSMA review, which presents both risk and opportunity for AI and digital sovereignty (Sections III, VI).

WINDOW FOR ACTION: The majority of new AI investment over the next half decade will reshape the technology stack, particularly for inference and deployment: the operational layer where AI systems process data, serve users, and generate value. Decisions being made today about infrastructure, platforms, and standards will shape the landscape for a generation, and making these decisions without regard for sovereignty risks locking Canada into dependencies that will be difficult to reverse. But the architecture is not yet settled, and a deliberate, sovereignty-oriented approach can influence how resources are deployed while the window for action remains open.

NAVIGATING TRADE-OFFS: Sovereignty does not come free. Sovereign infrastructure carries a cost premium, and the most secure options may constrain access to cutting-edge capabilities. Failing to adopt AI is itself a sovereignty risk: a country that falls behind becomes economically weaker, less competitive, and strategically vulnerable. This report presents a pragmatic, risk-based framework that calibrates measures to protect sovereignty with respect to their impact on economic competitiveness and the financial viability of the "sovereign solutions." The options presented are a menu for policymakers: different interventions can be pursued independently, at different speeds, and in different combinations.

CANADA'S STRENGTHS: Our layer-by-layer assessment of the AI technology stack (Section VII) reveals genuine strengths. Canada's clean energy advantage has attracted substantial data centre investment, and Canadian-owned operators provide some domestic alternatives to foreign hyperscalers at the cloud infrastructure level. Valuable government and private sector datasets in health, finance, law, and administration are strategic assets for sovereign AI development, particularly if Canada chooses to invest in protecting and developing them. Some Canadian AI application-layer companies have achieved global scale and multi-billion-dollar valuations. And a strong open-source ecosystem means sovereignty at the model operations layer derives primarily from infrastructure choices further down the stack.

CANADA'S VULNERABILITIES: The assessment also reveals critical weaknesses at the middle layers of the stack (Section VII). Cloud infrastructure is Canada's most acute controllable vulnerability: extraterritorial legal reach means that Canadian data residency does not equal Canadian data sovereignty, and recent incidents have demonstrated that foreign providers can deny service entirely under geopolitical pressure. The compute hardware layer presents equally severe vulnerabilities that cannot be addressed domestically. Most advanced AI chips and components are designed and fabricated outside Canada, with extreme concentration among a handful of foreign suppliers. At the foundation model layer, a small number of American companies dominate enterprise AI usage, and Canada has only one domestic alternative in Cohere.

Strategic Options for Canadian AI Sovereignty

1. Sovereign Cloud Infrastructure: Sovereign cloud infrastructure is the highest-priority strategic option presented in this paper. Cloud is Canada's most acute vulnerability, yet it is also the layer where realistic domestic solutions exist (Section VIII.B, Appendix B). Two primary models present distinct approaches to achieving cloud sovereignty:

››› *Juridical cloud sovereignty* (modelled on France and Germany) requires Canadian-owned and operated infrastructure that creates a legal air gap from foreign jurisdiction. This includes licensed operator arrangements where a Canadian entity operates foreign technology under Canadian control.

››› *Contractual cloud sovereignty* (modelled on Australia) achieves sovereignty through outcome-based controls — encryption keys held by Canadian entities, cleared personnel, and audit rights — that any provider can meet regardless of corporate nationality.

The appropriate model depends on the sensitivity of the workload and the institutional context governing it. One feasible set of minimum requirements could include self-hosted government infrastructure or juridical sovereignty for classified workloads; juridical sovereignty for sensitive personal and organizational data held by the government; and contractual sovereignty for private-sector data at the same level of sensitivity. Finally, for general business operations, which account for most cloud usage, market conditions should determine adoption.

For sovereign public cloud, pooling federal and provincial demand through arms-length Canadian sovereign compute providers, rather than fragmenting procurement across jurisdictions, would build critical mass for domestic providers and potentially narrow the sovereignty premium through economies of scale.

2. Additional AI Tech Stack Options: Strategic options extend beyond cloud to every layer of the AI tech stack (Section VIII):

››› *Compute hardware:* Canada could pursue a managed dependency strategy involving supply chain diversification across allied nations, bilateral assurance agreements linked to Canadian strengths in energy and critical minerals, multilateral semiconductor engagement, and contingency stockpiling for critical systems.

››› *Foundation model access:* Options could include establishing procurement preferences for Canadian providers, deploying open-source models on sovereign infrastructure as a strategic hedge, and diversification across model sources to avoid single-provider lock-in. Enterprise policies could address shadow AI, where most Canadian workers using AI rely on uncontrolled consumer tools rather than enterprise-grade alternatives.

››› *Data and data governance:* Options include introducing data governance frameworks that enable AI development while maintaining sovereignty and implementing measures to address Canadian content underrepresentation in global training datasets.

3. Cross-Cutting Enablers: Four policy mechanisms span multiple layers of the AI stack and require sustained attention to enable the strategic options outlined above (Section VIII.F). Procurement reform is needed to align government cloud and AI purchasing with sovereignty objectives, including updating the Government Cloud Framework with explicit sovereignty tiers and translating the Buy Canadian Policy into ICT-specific guidance. Workforce and security clearance requirements could be reassessed (including modernizing the security classification framework itself) to accommodate increased demand for qualified staff to operate sovereign infrastructure. Data portability requirements should ensure that switching between cloud providers remains practically feasible, not just theoretically possible. And domestic sovereignty audit capabilities are needed to validate provider claims on an ongoing basis, ensuring compliance rather than one-time certification.

4. CUSMA Preparation: The scheduled CUSMA review in July 2026 presents both risk and opportunity for Canada's digital sovereignty position (Section VIII.G). The United States has signalled aggressive positioning against what it characterizes as digital trade barriers, and U.S. industry groups are advocating for constraints on Canadian sovereignty measures. Canada should defend national security exceptions and government procurement carve-outs as non-negotiable foundations for sovereign cloud and AI policy. Digital sovereignty provisions should not become bargaining chips for concessions in unrelated sectors.

5. Foundation Model Training and Domestic Research: Even with robust sovereign AI inference capacity, Canada remains vulnerable if every model it uses originates from foreign suppliers or if model makers revoke access or restrict functionality. We outline two key options (Section IX). The first would be to support Cohere (one of very few foundation model companies outside the United States and China) more explicitly as a national champion, ensuring it has the resources and government support to remain competitive as global competition escalates. The second would be for Canada to pursue multinational frontier AI partnerships with like-minded democracies, pooling compute and committing to open-source collaboration, because no single middle power can sustain frontier model training alone. Canadian researchers also face severe capacity constraints that risk driving talent abroad, in which case domestic research compute would also require expansion.

6. Strengthening State Capacity and AI Leadership: Building sovereign AI requires matching state capacity (Section X). Canada's digital and AI governance is dispersed across at least six institutional actors, and no single entity possesses both the strategic authority and the operational capacity to drive a coherent sovereign AI strategy. To treat sovereign AI as a strategic national priority, the federal government should consider consolidating authority into a single institutional vehicle, including, potentially, a fully resourced Ministry of Digital and AI with cross-government delivery authority and/or a dedicated Sovereign AI Unit with investment capital and a clear mandate to buy, build, and invest in major sovereign AI projects. Institutional reform should be complemented by treating digital government modernization as a core component of sovereign AI strategy, and by establishing federal-provincial-territorial coordination on AI and data governance.

TIMELINES: AI sovereignty will not be achieved overnight, but meaningful progress is achievable by 2030 with deliberate action. The strategic options outline both immediate steps — data sensitivity tiers, procurement reform, security clearance modernization, CUSMA preparation, confronting shadow AI policies, and AI state capacity — and longer-term investments that can be made in sovereign data centres, multinational partnerships, and research infrastructure. Sovereignty is a spectrum, not a binary. Each action that reduces Canada's exposure to foreign leverage strengthens the country's position.

CANADA HAS THE INGREDIENTS: world-leading researchers, abundant clean energy, a world-class foundation model company, a growing sovereign infrastructure ecosystem, AI firms, and democratic institutions worth protecting. The window to act is open, but opportunities will narrow as global investment decisions harden into long-term commitments and advantages for other nations.

Part 2: Analytical Frameworks

The full report assesses Canada's AI position layer by layer through the AI technology stack, evaluating each layer against five dimensions of digital sovereignty. A third framework (the Data Sensitivity Spectrum) calibrates which sovereignty measures are appropriate for which workloads. Readers should consult the full report (Sections IV–VII) for the complete analysis; the summaries below provide the essential reference for interpreting the vulnerability assessment and policy recommendations that follow.

Digital and AI Sovereignty. Sovereignty — the ultimate authority to make and enforce binding rules without subordination to external authority — has never been absolute, but digital technology changes the character of the constraints. When a state cannot secure its data, enforce its laws over digital activity within its borders, or access computational resources necessary for economic competitiveness, its effective sovereignty is undermined. The full report identifies five dimensions of digital sovereignty (Section V), each concerning different objects and tested by different threats. AI amplifies these challenges: when AI models process government data through foreign-controlled infrastructure, they can compound every sovereignty dimension (jurisdictional exposure, operational risk, technological lock-in, societal distortion, and economic leverage) simultaneously.

Table V.1: Five Dimensions of Digital Sovereignty

Dimension	Core Question	Primary Threat
Jurisdictional	Can you enforce your rules?	Extraterritorial legal reach
Operational	Can you keep functioning under attack?	Cyber threats (attacks, espionage, exfiltration)
Technological	Can you migrate if needed? Can you substitute?	Vendor lock-in
Societal	Can you form and express preferences freely?	Democratic distortion, misinformation
Economic	Do you retain freedom to operate?	Economic coercion

The AI Technology Stack. The report analyses Canada's position across seven layers of the AI technology stack, ordered from foundational to applied: (1) **Data & Data Governance:** the information AI systems process and the rules governing its handling; (2) **Physical Infrastructure & Networks:** data centres, power systems, and network connectivity; (3) **Compute Hardware:** the specialized processors (GPUs, TPUs) that execute AI workloads; (4) **Cloud Infrastructure Services:** compute infrastructure offered as managed services by providers such as AWS, Azure, and Google Cloud; (5) **Foundation Models:** the large-scale AI models accessed for inference; (6) **Model Operations & Orchestration:** tooling to deploy, optimize, and manage AI models in production; and (7) **Applications:** end-user products and services built on AI capabilities. Sovereignty risks and dependencies differ at each layer; the vulnerability analysis and policy recommendations below are organized accordingly.

The Data Sensitivity Spectrum. Not all data requires the same protections. The report calibrates sovereignty measures to three tiers of data sensitivity, mapped across institutional contexts. The policy recommendations and minimum sovereignty levels referenced throughout this briefing note are anchored to these tiers.

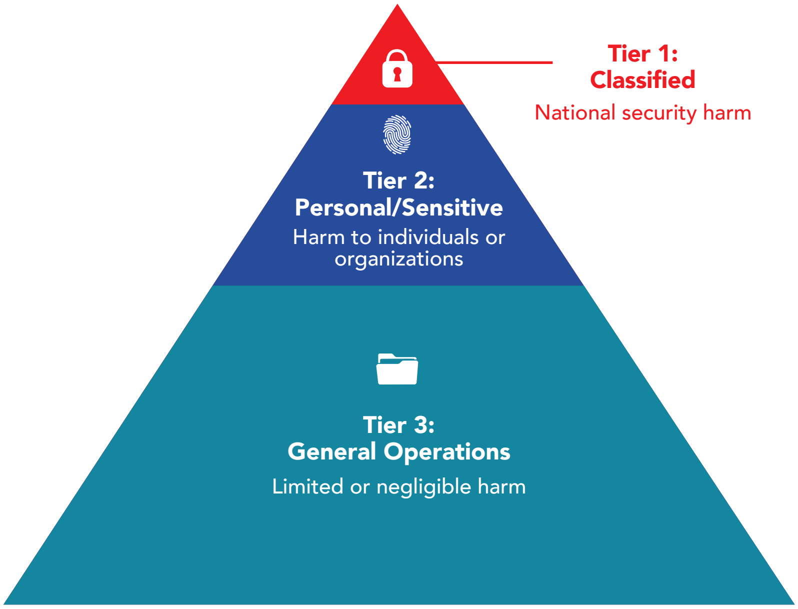

Table A.2: Data Sensitivity Spectrum

Tier	Harm Threshold	Equivalent Federal Classification	Federal Government	Provincial/Municipal Governments	Private Sector
Tier 1: Classified	Serious-to-grave harm to the **national interest** (military operations, intelligence sources, foreign relations, public safety)	Top Secret/Secret/ Confidential	Defence, intelligence, foreign affairs	*(Rare — shared classified data from the federal government)*	*(Rare — defence/ intelligence contractors with classified access)*
Tier 2: Personal/ Sensitive	Serious harm to **individuals or organizations** (identity theft, financial fraud, medical privacy violation, reputational damage, loss of competitive advantage)	Protected A/B/C	Citizen service delivery, personnel records, tax records, health admin	Health records, education, social services, vital statistics	PIPEDA-governed sensitive personal data (financial, health, biometric)
					Federally regulated sectors (banking, telecom, energy, transport)
Tier 3: General Operations	Limited or negligible harm (minor inconvenience, easily recoverable disruption, information already near-public)	Unclassified	Routine admin, open data, public communications, aggregate statistics	Municipal operations, open data, public notices	General business operations, consumer applications, PIPEDA-governed basic personal information (contact data, transaction records)

Part 3: Vulnerability Assessment Summary

The following chart maps the seven AI stack layers against the five dimensions of digital sovereignty, indicating both which dimensions are threatened at each layer and the severity of risk to Canadian AI sovereignty.

Table VII.2: Canadian AI Sovereignty Risk Heat Map

	Jurisdictional	Operational	Technological	Societal	Economic
1. Data & Data Governance	low	low	low	moderate	low
2. Physical Infrastructure	moderate	low	low	low	low
3. Computer Hardware	low	low	critical	low	critical
4. Cloud Infrastructure	critical	high	high	low	high
5. Foundation Models	low	low	low	moderate	moderate
6. Model Operations	low	low	moderate	low	low
7. Applications	moderate	low	low	moderate	moderate

LOW MODERATE HIGH CRITICAL

Table VII.3: Vulnerability Summary Across the AI Stack

Layer	Vulnerability	Key Sovereignty Threats
1. Data Governance	LOW-MODERATE	Canadian content underrepresented in training datasets; lacks mechanisms to leverage proprietary data assets for sovereign AI while maintaining privacy protections
2. Physical Infrastructure	LOW-MODERATE	Network routing through U.S. jurisdiction exposes Canadian traffic to foreign legal exposure; significant foreign ownership of data centre capacity
3. Compute Hardware	CRITICAL	Concentrated supply chain (NVIDIA dominance, TSMC fabrication monopoly); vulnerable to export controls and allocation decisions; no viable domestic alternatives
4. Cloud Infrastructure	CRITICAL	U.S. CLOUD Act jurisdiction despite Canadian data residency; service denial "kill switch" risk; vendor lock-in creates switching costs; concentration risk with hyperscalers
5. Foundation Models	MODERATE	Single-vendor concentration with Cohere as the only Canadian option; dependency on foreign providers creates revocation risk; American and/or Chinese political norms shape model values and behaviour; heavy consumer/employee usage creates lock-in beyond enterprise procurement
6. Model Operations	LOW-MODERATE	Hyperscaler MLOps platforms deepen cloud lock-in
7. Applications	MODERATE	"Shadow AI" creates uncontrolled data flows to foreign providers; consumer AI concentration (ChatGPT, Copilot, Gemini); Canadian AI-enabled companies strong, but underlying infrastructure choices matter

This assessment reveals vulnerabilities at every layer of the AI stack, but two stand out as requiring urgent attention. Cloud Infrastructure (Layer 4) represents Canada's most acute vulnerability with realistic domestic solutions. Compute Hardware (Layer 3) presents equally severe vulnerability but cannot be addressed through domestic action alone. Together, these critical middle layers concentrate the threats that flow through the entire stack.

Part 5: Policy Options Summary

The following table summarizes every policy option from the full report (Sections VIII, IX, and X), organized by theme / layer of the AI stack.

Theme/AI Stack Layer	Specific Policy Options for Consideration	Report Section
Sovereign Cloud Infrastructure	⋙ Establish three sovereignty models calibrated to workload sensitivity: self-hosted government infrastructure, juridical sovereignty (Canadian-owned operators with legal air gap from foreign jurisdiction), and contractual sovereignty (outcome-based controls any provider can meet) ⋙ Set minimum sovereignty levels by data sensitivity tier and institutional context (see Table VIII.1) ⋙ Aggregate federal and provincial demand (and/or from federally regulated sectors) through arms-length Canadian sovereign compute providers	VIII.B
Sovereign Cloud: Policy Mechanisms	⋙ Update Government Cloud Framework with explicit sovereignty tiers and evaluation criteria ⋙ Apply National Security Exception where appropriate ⋙ Create standardized certification frameworks for sovereign cloud providers ⋙ Provide direct support for domestic sovereign compute through procurement, financing, and potential equity contributions	VIII.B
Compute Hardware	⋙ Diversify supply chain across allied nations; cultivate multiple suppliers ⋙ Negotiate bilateral supply assurance agreements linking Canadian energy and critical minerals to technology access ⋙ Participate in multilateral semiconductor initiatives (e.g., Chip 4 Alliance) ⋙ Establish contingency stockpiles of GPUs and critical components for Tier 1 and Tier 2 systems	VIII.C
Foundation Model Access	⋙ Establish procurement preferences for Canadian model providers where possible; mandate Canadian models for Tier 2 and higher workloads ⋙ Deploy open-source and open-weight models on sovereign infrastructure as a strategic hedge ⋙ Require standardized interfaces, containerized deployment, and model portability across government AI contracts	VIII.D

Theme/AI Stack Layer	Specific Policy Options for Consideration	Report Section
Model Operations	» Adopt open-source MLOps and inference tooling on sovereign infrastructure » Procure from Canadian AI governance companies	VIII.D
Application Layer & Shadow AI	» Businesses should adopt enterprise AI policies requiring approved AI tools with data governance controls » Incentivize Canadian AI companies to deploy on sovereign infrastructure	VIII.D
Physical Infrastructure	» Mandate end-to-end encryption for sensitive government traffic traversing foreign infrastructure » Condition access to Canadian clean energy on sovereignty-enhancing commitments (data residency, local hiring, supply chain participation) » Support expansion of Canadian-owned data centre capacity	VIII.E
Data & Data Governance	» Modernize government privacy laws and data governance frameworks for AI-era requirements » Establish clear classification policies for government dataset use in AI training » Address Canadian content underrepresentation in global training datasets » Develop mechanisms to unlock strategic government datasets (health, financial, administrative, legal) for sovereign AI development	VIII.E
Procurement Reform	» Update Government Cloud Framework Agreements with explicit sovereignty tier requirements » Create procurement pathways for Canadian AI and cloud providers » Translate Buy Canadian Policy into ICT-specific guidance; consider categorizing AI infrastructure as dual-use » Align procurement requirements with Data Sensitivity Spectrum tiers	VIII.F
Workforce & Security Clearances	» Develop AI and cybersecurity talent pipeline » Scale clearance processing capacity to match sovereign infrastructure staffing demand » Modernize security classification framework using U.K. and Australian models	VIII.F
Data Portability	» Introduce data portability requirements for cloud providers » Mandate interoperability standards and demonstrated portability as cloud contract requirements » Require cloud-agnostic system design for government workloads	VIII.F

Theme/AI Stack Layer	Specific Policy Options for Consideration	Report Section
Sovereignty Audit	⋙ Develop domestic capacity to assess and validate provider sovereignty claims ⋙ Create standardized assessment frameworks leveraging work of Digital Governance Council ⋙ Establish ongoing validation requirements, not just initial certification	VIII.F
CUSMA Review Preparation	⋙ Defend National Security Exception and government procurement carve-outs as non-negotiable ⋙ Build coalition with like-minded partners (Australia, U.K., Mexico, EU) ⋙ Maintain space for outcome-based, origin-neutral sovereignty requirements ⋙ Refuse to trade digital sovereignty provisions for concessions in unrelated sectors	VIII.G
Foundation Model Training: Cohere	⋙ Provide R&D funding and talent retention incentives ⋙ Consider treating Cohere as a "national champion" given its rare status as a commercial foundation model provider; support adoption of Cohere models across government and Canadian businesses	IX.B
Multinational Frontier AI Partnership	⋙ Pool sovereign compute resources with partners (EU member states, U.K., Australia, Japan, South Korea) ⋙ Commit to producing fully open-source foundation models as a key partnership output ⋙ Build on the Abecassis et al. (2025) "Blueprint for Multinational Advanced AI Development" framework	IX.C
Research Compute	⋙ Expand Canadian dedicated research compute capacity ⋙ Link research compute investment to multinational partnership participation ⋙ Strengthen talent retention through competitive research infrastructure ⋙ Design research infrastructure for potential dual-use applications	IX.D
State Capacity	⋙ Increase state capacity for AI sovereignty by: 1) Creating a fully resourced Ministry of Digital and AI with a cross-government, delivery-focused executive empowered with policy authority and a CIO with operational authority, and/or; 2) Establishing a dedicated Sovereign AI Unit with investment capital and a mandate to buy, build, and invest in sovereign AI projects	X.B

Theme/AI Stack Layer	Specific Policy Options for Consideration	Report Section
Enabling Conditions	⋙ Anchor digital government modernization within the same institutional home as AI strategy ⋙ Establish a federal-provincial-territorial working group on AI and data governance to coordinate procurement standards, shared infrastructure investment, regulatory alignment, and Indigenous data sovereignty principles	X.C

Compute Infrastructure: Determining Minimum Recommended Sovereignty Levels

Table VIII.1 presents one feasible set of minimum sovereignty levels to each combination of data sensitivity tier and institutional context. The assignment reflects both the risk profile of the workloads and the policy levers available in each institutional context.

Table VIII.1: Minimum Sovereignty by Tier and Institutional Context

Tier	Federal Government	Provincial/ Municipal	Federally Regulated Industries & CCSPA Sectors	General Private Sector
Tier 1 (Classified)	Self-hosted or Juridical	N/A	N/A	N/A
Tier 2 (Personal/ Sensitive)	Juridical	Juridical	Contractual	Contractual
Tier 3 (General Operations)	Contractual*	Contractual*	Marketplace Determined	Marketplace Determined

**In practice, there is potential to pool government demand across tiers and/or levels of government. See discussion on pooling below.*

The table identifies the lowest standard sufficient for each combination of sensitivity and institutional context; organizations are always free to exceed the minimum. For **Tier 1** (classified) workloads, self-hosted infrastructure is most appropriate for Canadian Eyes Only data, while juridical sovereignty suits systems requiring interoperability with allied intelligence partnerships. Tier 1 workloads are almost exclusively federal. For **Tier 2** (personal/sensitive) data, juridical sovereignty is the recommended minimum for both federal and provincial/ municipal government, reflecting higher public expectations of government data stewardship. For the private sector, the recommended minimum is contractual rather than juridical — the private sector is too fragmented for a uniform juridical mandate — but a contractual floor remains important, as the harm from a breach does not diminish because of who holds the data.

For **Tier 3** (general operations), the recommended minimum for government workloads is contractual sovereignty, ensuring a baseline of Canadian-held encryption keys and cleared personnel even for lower-sensitivity applications. For the private sector, adoption at this tier should be marketplace-determined, though many enterprises may voluntarily adopt contractual or juridical options for competitive, reputational, or risk-management reasons.