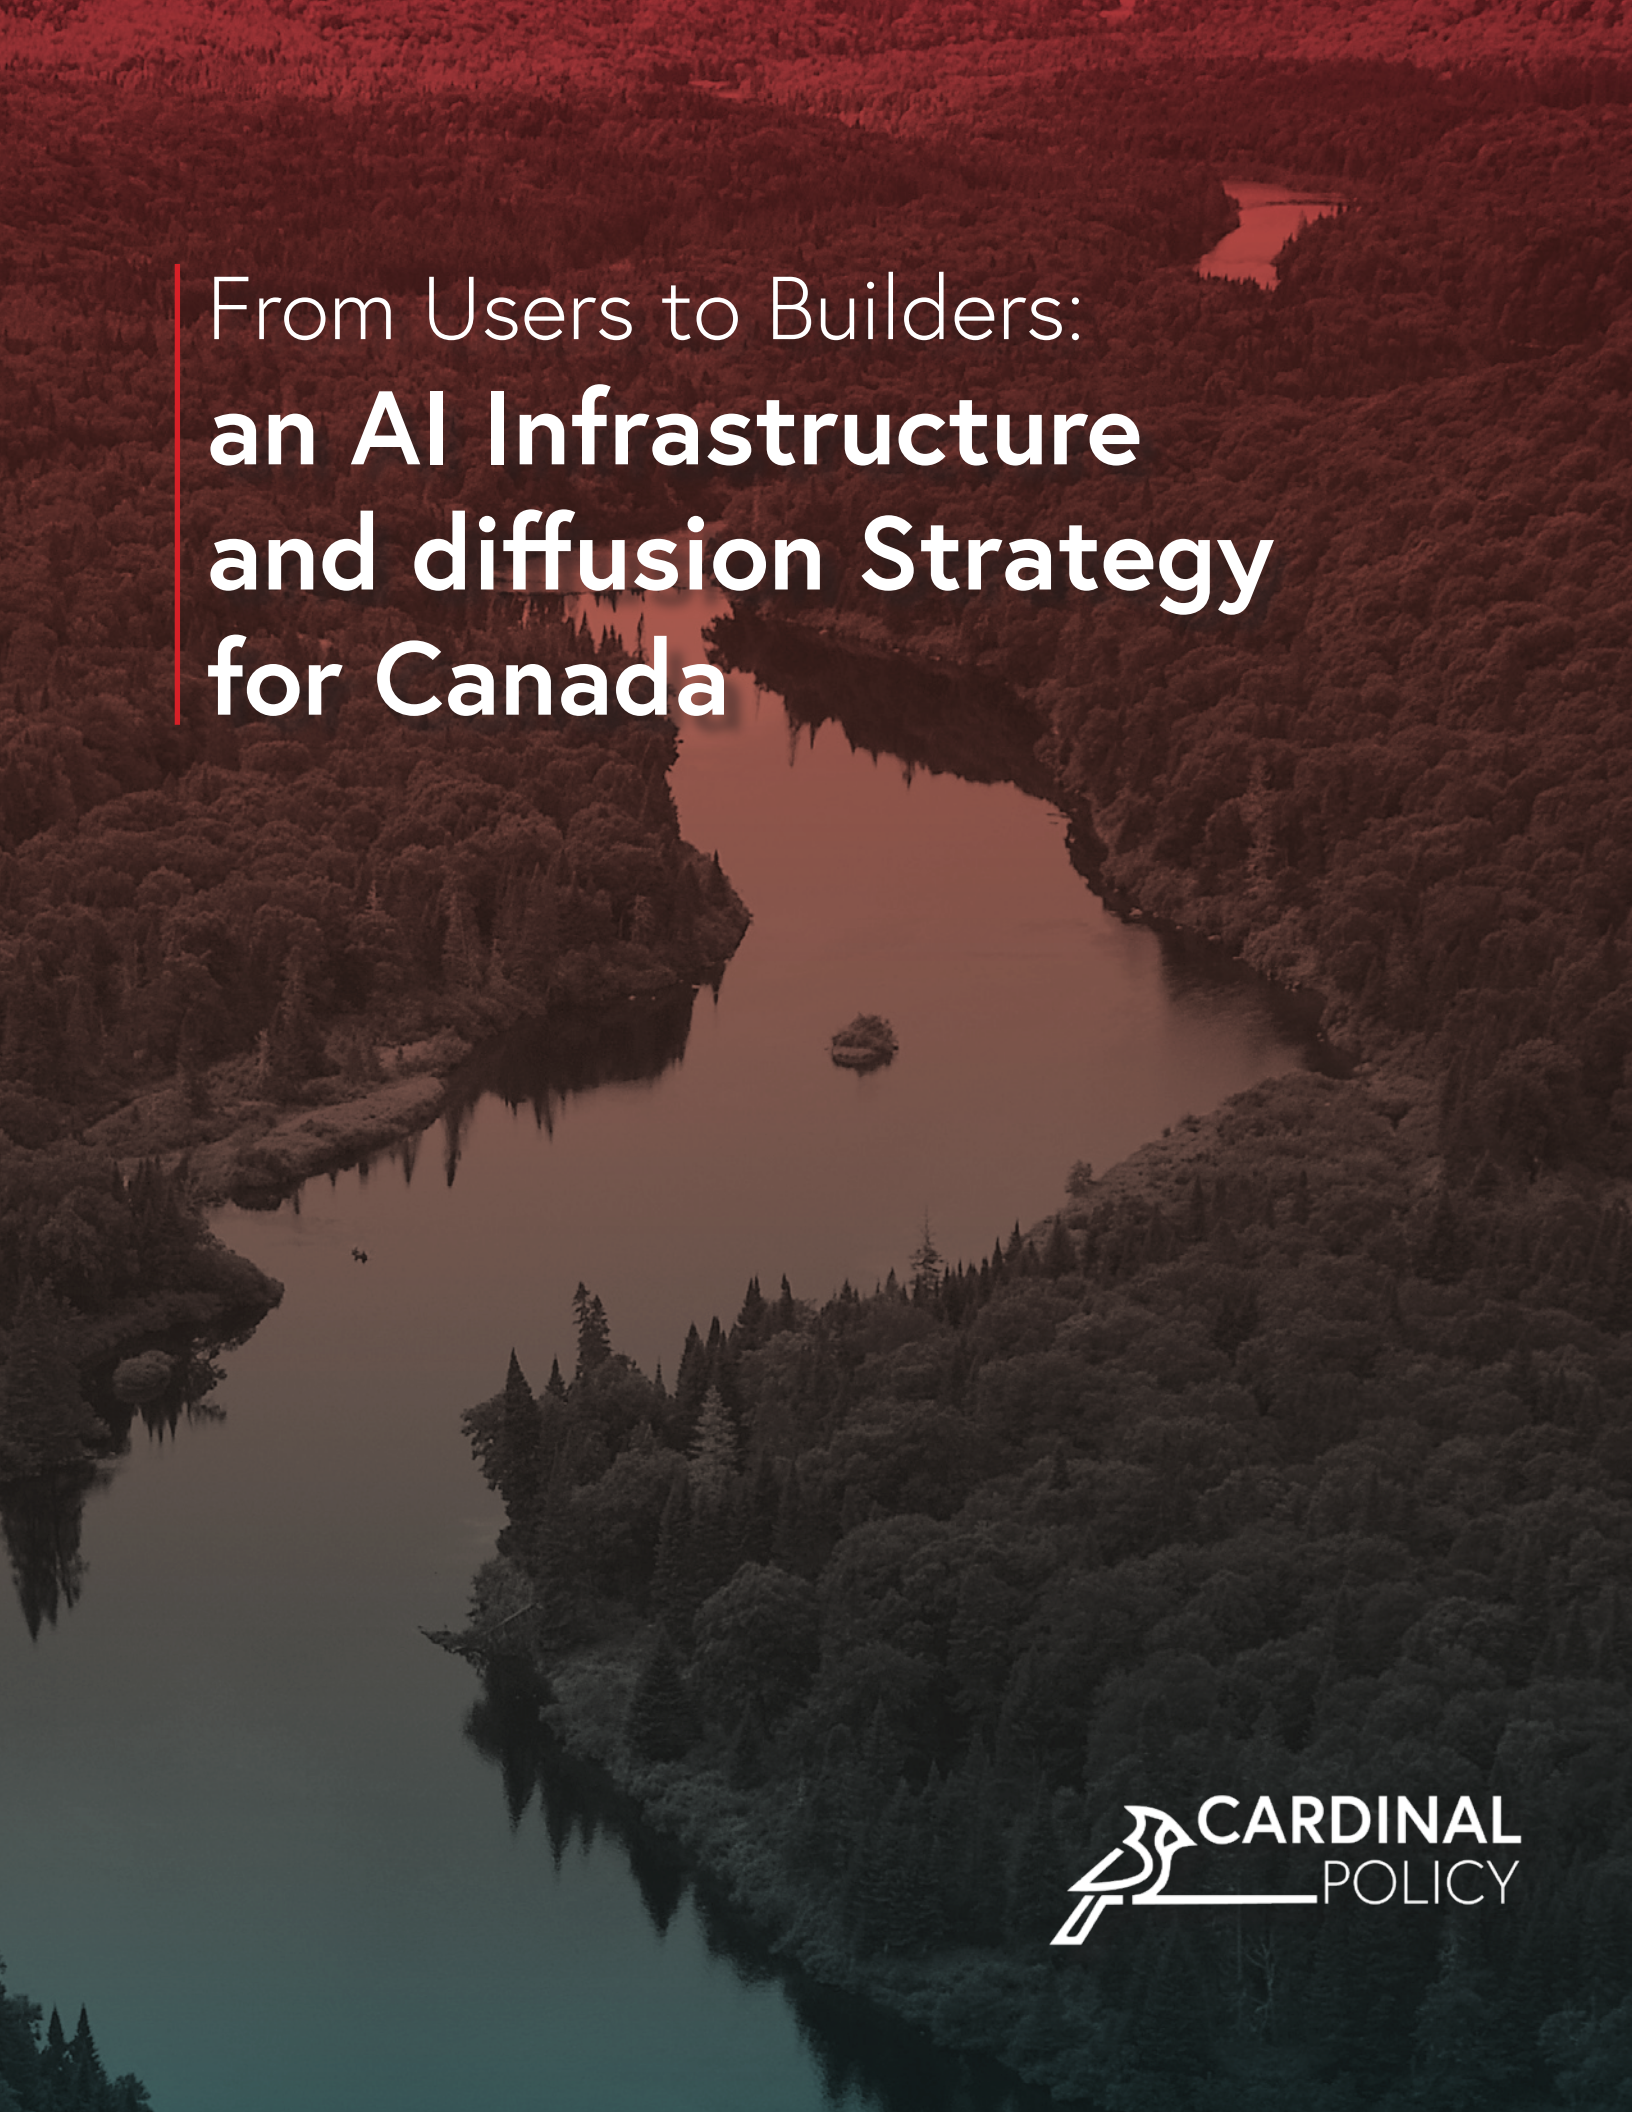

From Users to Builders: **an AI Infrastructure and diffusion Strategy for Canada**

CARDINAL POLICY

AUTHOR'S NOTE

This paper was prepared as a submission to the Minister in the context of my role on the Canadian AI Strategy Task Force. I was specifically asked to examine AI infrastructure and sovereignty, which explains their preeminent focus in this paper. However, given my experience with the AI Adoption Initiative (AIAI) over the past year and my research with the Bennett Institute for Public Policy at the University of Cambridge, I felt compelled to also address the broader challenge of AI adoption and productivity growth in Canada.

This document is intended not as an academic exercise but as an action-oriented set of proposals for industrial policy. While many of the views expressed are my own, I have sought to ground them in empirical analysis and relevant research wherever possible.

I am grateful to Minister Solomon for the opportunity to contribute to Canada's AI ambitions and to serve Canadians in this capacity. I would also like to thank the hundreds of individuals who shared their insights through emails, calls, coffees, lunches, and roundtables. Their perspectives deeply informed my understanding of the Canadian AI ecosystem, and many of their arguments and ideas are reflected throughout this paper.

Finally, AI policy and geopolitics evolve rapidly. This paper represents my best analytical effort as of October 2025. A revised version in 2026 would no doubt reflect new facts and perspectives. As John Maynard Keynes famously said, "When the facts change, I change my mind."

Marc-Étienne Ouimette
President and Founder, Cardinal Policy

ABOUT THE AUTHOR:

Marc-Etienne Ouimette
President and Founder, Cardinal Policy

Marc Etienne Ouimette has over 15 years of experience in policy and deep tech. He is the founder of Cardinal Policy, where he advises firms and governments on the intersection of AI, public policy, and industrial strategy. He is also an affiliate researcher at the University of Cambridge's Bennett School of Public Policy where he does research on AI sovereignty and industrial policy.

As the head of global AI policy at AWS, Marc worked closely with senior leadership to develop strategic policy positioning for the company and led engagements with numerous governments and multilateral bodies. He notably sat on the B7's advisory committee, was a member of the OECD network of experts, and was board chair of the Global Partnership on AI's Centre of Expertise. Prior to AWS, Marc led the global policy team at Element AI from its founding until its acquisition by Service Now, was general counsel at Mitacs, and was a corporate lawyer at one of Canada's major firms. In October 2025 he was named by the Canadian Minister of Artificial Intelligence to the Canadian AI taskforce responsible for developing the country's new AI strategy. He also supports youth education in AI through his role as board chair of Digital Moment, a national organization dedicated to training the next generation in AI & the digital space.

ABOUT CARDINAL POLICY:

Cardinal Policy is a geopolitical and Canadian policy consulting firm focused on artificial intelligence and its associated technology stack. We help the world's leading organizations understand, navigate, and respond to technology policy issues that will determine how industries evolve, economies grow, and nations compete in the 21st century. Cardinal Policy works in strategic partnership with Minerva Policy Futures.

TABLE OF CONTENTS

Introduction 5

Resilient AI compute capacity 6

The Canadian Compute Innovation Initiative 7

The Strategic Adoption Cluster 9

Maximizing CSP Sovereign controls 11

The Copyright Issue 12

Government's capacity to execute 12

Accelerate AI diffusion 14

Priority sectors 16

Industry Digital Plans 16

Ai adoption Talent Pipelines 17

Tax Policy 18

Public sector reference deployments 19

Leverage standards 19

Data readiness and access 20

Open-source and the difussion fly-wheel 21

Cultivating demand and culture 22

Conclusion 22

Annex 1 - Sovereign Cloud Commitments 24

Annex 2- Data & Compute Embassies 25

Annex 3- Copyright & AI 26

References 28

01 INTRODUCTION

Artificial Intelligence (AI) is a general-purpose technology (GPT), akin to electricity or the internet. Despite continuing research to advance the frontier of AI, AI is well established in our everyday lives, advancing in commercialization to the point that we often don't recognize it as "AI" anymore (Google Maps directions, Netflix recommendations, etc.). As such, trying to craft one broad "AI Policy" or strategy at this stage of the technology is as futile as having one policy for "electricity", "engines", or "computers". For an AI strategy to be of any use, it must be broken down into components that reflect the technology's maturity, the complexity of the AI "stack", the objectives the country is trying to achieve, and the capacity of that country (in research, natural resources, existing leading enterprises, in its diplomatic links, and so on) to execute against the objectives.

Things were different at the time of the first national AI strategy put out by Canada, in 2017. AI was still very early in its technological development. As such, a strategy broadly based on securing fundamental research capacity through talent (CIFAR grants, specific visa streams, creation of research institutes, etc.) made sense.

Today, a national AI strategy with the primary tactic of securing fundamental research talent would be limiting itself to a sliver of AI's economic and geopolitical impact. A 2025-26 strategy needs to properly account not only for the evolving AI tech stack but for geopolitical tensions and supply chain constraints. It must understand the levers at the government's disposal and its capacity to execute. And as we're dealing with multi-billion-dollar investments, it must struggle with details while avoiding decision-making paralysis.

In September 2025, Prime Minister Carney announced his desire to build a "Canadian Sovereign Cloud". As he put it, "This would build compute capacity and data centres that we need to underpin Canada's competitiveness, to protect our security, and to boost our independence and sovereignty (…) This will give Canada independent control over advanced computing power while reinforcing our leadership in AI."[i] This is an understandable statement from the Prime Minister at a moment when dependency on US partners and infrastructure is reassessed. A "Canadian Sovereign Cloud" may be a laudable north star. But, as I lay out below, as a strategic objective a fully sovereign Canadian cloud is unattainable in the short or medium term, even without contemplating chips. Over the long-term, a sovereign compute strategy is better framed as an ethos, one that guides the Canadian strategy towards digital infrastructure that is resilient to supply chain and geopolitical pressures and that at its core supports Canadian economic and productivity growth.

Against this backdrop, I argue that Canada's 2025–26 AI strategy must focus on two distinct but mutually reinforcing objectives: 1) building resilient compute capacity, both by partnering with CSPs and model developers (the SAC model described below), and by developing a comprehensive program to build up domestic AI compute capabilities (the proposed Canadian Compute Innovation Initiative, below) ; and 2) drive broad adoption of AI across sectors to boost productivity, create domestic champions, and increase Canada's techno-geopolitical relevance. Rather than aim for symbolic autonomy, it should pursue a "third way" of Cooperative AI sovereignty, identifying where Canada can credibly lead, and where it must partner. The following paper offers a pragmatic blueprint for doing so.

Taken together, these actions would give Canada real bargaining power, national resilience, and economic lift. Crucially, they also align with the historical lesson captured in Jeff Ding's work: national power through general-purpose technologies accrues less to the first inventors than to those who work to diffuse them widely.

02 RESILIENT AI COMPUTE CAPACITY

Digital sovereignty concerns are not new. Their interpretation in the AI space is a more recent affair, however, pushed to the front of the stage by several factors: advances in AI and the release of ChatGPT; a rise in concerns relating to the resilience of global supply chains (as exposed by COVID and, more recently, Chinese moves to restrict access to rare earths and other critical minerals used in key technology components); concerns about technology dependencies (including dependencies on Taiwan), as outlined in Chris Miller's "Chip War", Edward Fishman's "Choke Points", and Abe Newman and Henry Farrell's "Underground Empire"; a global PR/GR campaign by Nvidia; and, most recently, the second Trump administration trade wars, and accompanying concern that the US will leverage dependencies in trade negotiations (or other settings).

Consequently, many countries have developed "Sovereign AI" plans aimed at securing a degree of independence and control over the AI compute stack. Countries have used different tactics at different levels of the stack, ranging from model development, to financial support for key domestic players, to striving for a degree of "ownership" over the compute capacity required to train and operate AI.

Canada's goal should be to ensure it can withstand external pressures to its compute capacity while remaining deeply engaged in allied AI networks. Amid intensifying U.S.–China competition, Canada will face pressure to align with the US while simultaneously guarding against overdependence.

Canada and other AI middle powers cannot aim for "AI dominance", as the US government puts it. Instead, they must strive for a resilient AI supply chain that can withstand a high degree of political and supply chain pressure. Canada has wisely realized it cannot currently hope to fully replicate a CSP offering through public means, and it cannot replace every component of the AI value chain with Canadian (or even non-US, non-Chinese) private sector partners. But, by breaking down tactics along the components of the AI Compute supply chain, having a crisper understanding of the type and size of compute capacity required, and by assessing potential bilateral or "mini lateral" partnerships with like-minded countries to pool funding and leverage each other's capacities, Canada can meaningfully diminish its exposure, increase its technical and geopolitical relevance, and stimulate domestic productivity.

COMPUTE CAPACITY

The compute infrastructure strategy I suggest pursues three parallel paths:

1) Canada should have a long-term plan to support the development of domestic capabilities along the entire AI compute stack by reconceptualizing the Sovereign AI Compute Program as the Canadian Compute Innovation Initiative, with the objective of increasing the country's agency in compute capacity over time ;

2) Canada should leverage the demand of private sector partners for energy supply and compute (including CSPs and frontier model developers) to incentivize the development of a Strategic adoption cluster, should explore Joint Venture (JV) models to achieve this, and should use its position to ensure a non-insignificant percentage of the compute capacity within this build-out be reserved for domestic AI adoption or for open-source models development;

3) Canada should negotiate the highest degree of sovereign controls offered by CSPs for critical workloads in CSP build-outs.

It is not the objective of these AI compute projects to match frontier model compute build-outs in scale. Canada should neither concede the frontier entirely nor chase an unrealistic goal of matching the U.S. or China in frontier compute capacity. Beyond a certain point, additional training compute capacity yields vastly diminishing returns for national needs, unless one is pursuing true frontier capacity[III]. The objective is to have a controlling of compute capacity is to have a controlling stake in resilient capacity *in country*, not to beat Google or OpenAI on benchmarks. This strategy also avoids over committing to scaling laws as the sole basis of LLM capacity growth[IV].

1 Build up domestic AI compute Capabilities through the Canadian Compute Innovation Initiative (CCII)

Some Canadian compute demand will be met through the national supercomputing facility to be built by ISED as part of the Sovereign Compute Program. While welcomed, this Program is mostly aimed at a sophisticated research audience. I suggest Canada build on this program to launch a Canadian Compute Innovation Initiative (CCII), a long-term program aimed at building domestic capacity up the AI compute stack. While it could immediately service sophisticated users, it should be envisioned as a domestic advanced TRL version of "Bell Labs" for compute, attractive to researchers striving for breakthroughs while simultaneously delivering capacity across the stack.

By focusing first on users able to leverage bare-metal or lightly orchestrated compute environments (the relatively simplest components of the stack to assemble) Canada can rely on domestic providers to anchor the earliest stages of investment. Institutions such as MILA, the Vector Institute, and AMII, as well as certain university researchers will be primary beneficiaries, able to train, test, and fine-tune models. Select Canadian AI companies with similar compute know-how and capabilities could also access this facility for subsidized model development. Together, these can anchor a *sovereign-compute* core, while forming a launch platform for a broader ecosystem whose aim, over time, should be to service a broader spectrum of users.

For other less technical users (most start-ups, SMEs, and traditional enterprises starting their AI journey) such bare-metal or lightly orchestrated compute resources require too much technical expertise to be useful. This is one reason why many if not most Canadian users will continue to require hyperscaler-class (CSP) offerings in the near and medium term. Preventing domestic users from accessing this capacity through CSPs to prioritize domestic adoption would be detrimental to Canada's ability to adopt AI rapidly and at scale. Instead, Canada should focus on ensuring that its domestic compute capabilities program can interoperate with CSPs, while having a clear plan to capture more of the compute value chain over time.

The AI compute stack is complex and multilayered. At the core lies energy supply, energy-efficient compute and networking hardware such as GPUs, CPUs, and high-speed interconnects; above that are the middle layers of orchestration and resource management systems, including Kubernetes, VMware, or Slurm, along with developer platforms such as PyTorch and TensorFlow, and MLOps tools like Kubeflow, MLflow, and KServe that manage model training and deployment; while the upper layers consist of the human interface, e.g. solutions architecture, user experience design, and customer support that make infrastructure usable, reliable, and attractive to users.

Canada already has varying degrees of capacity at different layers, and arguably full capacity at the bare metal compute layer, in partnership with chip providers (Hypertec is an OEM partner of Nvidia, for instance). Over time, the CCII should aim to grow outward from this core "full" bare metal capacity, much as CoreWeave did by transforming itself from a "scrappy cryptominer" to a much-nearer to full stack offering leveraging open-source tooling[V]. Unlike most former crypto-mining operators that remained hardware-centric, CoreWeave moved up the software stack, emphasizing open-source orchestration, HPC scheduling, and developer-platform services.[VI] A company like Calgary-based Denvr Dataworks is clearly aspiring to follow this path and position themselves

as an AI cloud provider, intending to wrap orchestration, developer UX and platform tools around their infrastructure. The CCII would provide an opportunity for such firms to collaborate on applied R&D while simultaneously providing compute capacity to sophisticated domestic AI customers.

The AI Sovereign Compute Infrastructure Program should be reconceptualized as a long-term program of investment not only in hardware but in capacity development in software and talent, to attract and train an expanding pool of systems engineers, solutions architects, and UX designers who can translate user needs into high-performance, multi-tenant environments. The program would provide a much-needed stimulus and forcing function to drive partnership between the Canadian innovation sector, academia, and the public sector, with a mission specific objective: supporting the next generation of Canadian compute capacity and resiliency, by aligning capabilities and roadmaps for the future.

The Government could accelerate this by building out the (current) one-time ISED Sovereign-Compute Program investment into a 10-year program operating in partnership with Canadian businesses in the AI compute space. Financial support should be tied to expanding the depth and breadth of the pool of users, and leverage existing finance mechanisms where possible (Canadian Infrastructure Bank, NRC, etc.). The CCII should also work with open-source foundations to support open-source projects by domestic firms aimed at closing functional gaps of stack components.

The CCII should partner with MILA and/or universities to actively recruit and embed a cohort of expert systems engineers and developer-platform specialists drawn from academia and industry. These individuals would bring deep experience in high-performance computing, distributed systems, container orchestration and AI workflow architecture. By seconding such experts into the CCII, the government would be investing not just in hardware, but in human capability. Over time, these specialists can act as knowledge-multipliers, guiding training programs and mentoring newer engineers in domestic providers. The point is to recruit people who understand both "bare metal up" and "developer experience down" to close the gap between raw compute and usable AI platforms.

CoreWeave's rapid rise was driven largely by engineers and executives recruited from commercial CSPs[VII] (especially Google Cloud and AWS), not academia. Those individuals were motivated by equity, startup dynamics, and private-sector compensation, which makes them unlikely to move into a public supercomputing organization or university. Recruiting such experts will need to happen in partnership with for-profit domestic providers. The Government should explore forms of support (such as joint appointments, or industrial research chairs with MILA or other academic-adjacent organizations) to help with their recruitment.

Another way to accelerate domestic upskilling would be to structure part of the CCII through a call-option build-operate model with a proven AI cloud actor. Under this approach, an experienced AI-cloud provider would co-develop and operate an advanced facility in partnership with the CCII and Canadian infrastructure providers. The model would combine private-sector execution expertise with public oversight and knowledge transfer. Over a defined period (typically three to five years) the partner would run the facility while training Canadian engineers, systems architects, and operations teams in large-scale AI-cloud management. The Government would retain an option to assume full or partial ownership once local capacity has matured, ensuring that the technical know-how and operational playbooks remain in Canada. This approach would de-risk the early phase of deployment, while embedding the practical, hands-on experience required to sustain sovereign capability over the long term.

Canada should also **invest in compute efficiency** to **chip away at the differentiating capacity of brute-force compute**. Funding research on techniques like distributed and low-bandwidth training, model compression, and novel AI chip architectures can yield order-of-magnitude improvements that reduce the gap between small and big players; a dollar spent on efficiency research might save ten in future infrastructure costs. The proposed MILA/Université de Sherbrooke/C2MI research centre would be a step in the right direction and should be integrated within the CCII.

2 Explore Joint Venture (JV) vehicles and financing mechanisms to support the build-out of a "Strategic Adoption Cluster"

Canada should explore the feasibility of a joint-venture infrastructure platform that brings together government, energy providers, technology OEMs, and AI model developers under a common framework. Rather than a single government-owned facility, it should be conceived as a hybrid public–private investment vehicle that pools capital, aggregates demand and ensures sovereign access to advanced compute capacity.

This approach would leverage the fact that AI compute infrastructure has become a form of energy-linked digital infrastructure whose economic and strategic value depends not only on the compute hardware but also on the energy systems and networking that sustain it. These arrangements for massive investments in AI capacity are possible because *training* AI models, including frontier models, does not run into the same latency and proximity considerations as data centres built largely for *inference*. While CSPs have historically preferred to locate AI infrastructure in the same regions as their other data centres, energy constraints and high demand for training capacity has opened new, often remote geographies for data centre buildouts (e.g. the Gulf States).

As such, a key component of being a geopolitically relevant middle AI power for AI compute (and a strong negotiation stance in US "tech deals") might not be 'a state that has a lot of compute', but 'a state to whom you can send an H100 that can actually get it up and running'.[VIII] Canada is not without assets on this front. The country boasts plentiful clean energy potential for data centres, a cold climate, large institutional investors capable of multi-billion co-investments[IX], telecom giants with infrastructure capabilities, and strong research capacity.

Positioning a strategic adoption cluster (SAC) as an asset tied to clean energy generation would allow it to attract institutional investors such as pension funds, infrastructure banks, and sovereign wealth funds. These actors are already experienced in financing long-duration infrastructure assets and would likely view such an initiative as a logical extension of their portfolios into digital infrastructure.[X]

The proposed SAC joint venture would involve four types of partners: foundational model developers, energy and infrastructure operators, technology manufacturers, and public institutions. Model developers such as Cohere, Anthropic, Mistral and Open AI would act as anchor tenants by committing to use a portion of the facility's compute capacity through long-term contracts (guaranteed "off take"). Domestic energy and infrastructure operators would provide power, site development, and systems integration. Technology manufacturers such as NVIDIA or AMD would supply hardware and integration expertise.

Public participation should focus on guaranteeing domestic access to a portion of the compute capacity and guiding strategy rather than on full ownership. This can be achieved through callable equity and structured capacity reservations that secure long-term national access. To make the facility financially viable and attractive to private credit markets, several instruments can be incorporated. The Government should focus on establishing long-term renewable energy contracts, or Power Purchase Guarantees, to ensure stable operating costs.

Agreements with hardware manufacturers could include GPU buyback clauses, allowing equipment to be repurchased at a pre-determined residual value after several years, which would reduce depreciation risk. Sovereign loan guarantees would be applied only to the power, cooling, and interconnect components, which are durable and (Canadian) low-volatility assets, while excluding high-depreciation compute hardware. Together, these measures would likely make the SAC a bankable and sustainable investment while minimizing public exposure.

The governance structure must embed sovereignty and oversight without introducing unnecessary amounts of bureaucracy. A Public Technical Trustee representing the Federal government should be embedded within the governance framework. This Trustee would have specific authority over access and the ability to block decisions that compromise sovereign control or operational continuity. Additionally, the joint venture should include a Right of First Refusal clause, allowing the Government to participate in future expansions or secure additional capacity before external partners are invited. These provisions would ensure long-term alignment between national interests and commercial expansion.

Most of the facility's capacity would be secured through long-term off-take agreements with one or more anchor tenants, such as OpenAI, Anthropic, Cohere, or other major model developers. These anchor commitments would function as "take-or-pay" GPU-hour contracts, guaranteeing baseline utilization and enabling the project to attract private financing at infrastructure-grade terms. The remaining capacity could then be made available to Canadian or allied institutions or critical partners. This balance would preserve the SAC as both a sovereign resource and a commercially investable facility.

To further align the SAC with national AI-adoption goals, Canada could use its position in the joint-venture vehicle to introduce a conditional allocation mechanism under which model-providers receive access to additional compute capacity only if they demonstrate specified commitments to domestic uptake. For example, future capacity tranches might be granted on the basis that a defined share (say 70%) of the corresponding capacity is tied to new Canadian clients or applications via service or licensing agreements. In this way the Government leverages the SAC both as a commercial platform and as an industrial policy instrument: it encourages model-providers to compete for allocation by expanding Canadian adopters, thereby creating a flywheel of compute-access → domestic AI uptake → stronger justification for further SAC expansion → further allocation. Embedding such an incentive structure within the JV's governance and allocation framework transforms the SAC from a static sovereign asset into a proactive ecosystem lever.

Any required subsidization would occur at the level of Canadian adopters (through a targeted voucher program redeemable exclusively for SAC-based compute) rather than at the model-provider level. In China, several municipal and provincial governments have introduced "computing power vouchers" that subsidize the rental of AI inference for SMEs. These vouchers, in some cases covering up to 80 percent of rental fees[XII], directly lower the barrier to entry for SMEs seeking to adopt AI tools.This approach would simultaneously stimulate domestic demand and create a qualified pool of prospective clients, enhancing the SAC's attractiveness for model providers and ultimately as a non-speculative, infrastructure-grade investment vehicle.

Additionally, to complement long-term off-take arrangements, a portion of the SAC's capacity could be allocated through a flexible, futures-style mechanism analogous to proposals in the U.S. AI Action Plan[XIII]. Under this approach, smaller increments of compute would be auctioned or contracted on shorter time-horizons, enabling start-ups and SMEs to access advanced hardware without long-term commitments. The mechanism could be designed to reinforce the domestic-adoption flywheel by linking eligibility or allocation priority to demonstrable Canadian use-cases. Model providers that channel capacity through such domestic service agreements would strengthen their case for subsequent allocations under the main SAC framework.

Over time, the SAC could evolve into the nucleus of a broader allied initiative that could also leverage the Data and

Compute Embassy model **(Annex 2)**. This model would allow Canada to pool demand from trusted partners under common governance and technical standards. Over the long term, the SAC could position Canada as a reliable hub for democratic AI infrastructure, providing both secure capacity and a model for international cooperation.

This model would deliver several important outcomes. It would secure sovereign compute capacity for Canada and its partners while increasing domestic control over critical compute infrastructure. It would attract long-term private capital through predictable, energy-linked returns. It would likely build domestic expertise in high-performance computing, orchestration, and systems engineering alongside the CCII. Finally, it would position Canada as both a sovereign user and a strategic supplier of trusted AI compute capacity.

In summary, the SAC should be developed as an investable, energy-backed infrastructure platform that blends public oversight with private expertise. Its success will depend on strong governance, clear demand commitments, and long-term partnerships with both domestic and allied stakeholders. Over time, the SAC can evolve into a cornerstone of Canada's sovereign AI strategy, combining national resilience with global integration.

3 Maximize the sovereign controls offered by CSPs

Physical location of compute capacity still matters. Despite the Westphalian system's limitations, governments maintain supreme authority to govern within their defined physical territory. This allows them to regulate physical compute infrastructures located in their country in a way that other governments cannot, notably to enforce "Know-Your-Customer checks, audits, or other controls at the point at which AI models and data sets enter a data centre for training or deployment".[XIV] This does not entail full sovereign control. The nationality principle allows a state to assert jurisdiction over its nationals, whether individuals or legal entities, even when their conduct occurs abroad. Under the principle of nationality, data center operators can serve as jurisdictional "hooks" into data centers abroad."[XVI] Moreover, the U.S. Foreign Direct Product rule extends U.S. export control jurisdiction extraterritorially by subjecting foreign produced items (such as servers, chips, or cloud infrastructure) if they are the "direct product" of U.S.- origin-technology, software, or equipment, regardless of where they were manufactured[XVII]. Still, when push comes to shove the physical location of the infrastructure ultimately allows physical control over it if necessary, while several different legal structures can offer additional levels of control to increase the likelihood of continued operational autonomy.

The reality is that partnering with CSPs to meet domestic AI compute infrastructure demand will be necessary for the foreseeable future, including for critical workloads, and that many if not most Canadian companies and government entities rely on their capacity and full stack know-how for their AI journey[XVII]. CSPs also offer multi-region resilience, international presence, elasticity to meet peak demand, and more. For these reasons and others, CSPs will be part of the solution in the short, medium, and likely long term.

Still, CSPs have digested that there are rising sovereignty concerns amongst key customer nations. Despite pushing back on the legitimacy of said concerns and accompanying regulatory and procurement controls, they have driven into the "sovereignty market", as Pablo Chavez has well documented[XVIII], offering governments various means to ensure operational autonomy should the CSP be required to cease its services in-country or region (see **Annex 1**).

Canada should explore how it can maximize these controls when partnering with CSPs for critical workloads. For example, those supporting essential government functions, sensitive research, or applications with national security implications. Limiting their scope in this way avoids imposing unnecessary cost and compliance burdens on general compute demand while ensuring that the workloads most vital to the public interest are better protected.

At the same time, the quantity of these workloads must be sufficient to justify the additional investment by CSPs. Canada could draw on the data- and compute-embassy model described in Annex 2, in which a small number of sovereign countries aggregate and host critical workloads under shared technical and legal frameworks. Pooling protected demand across like-minded nations could create the scale and predictability of usage needed to warrant deeper CSP investment in sovereign infrastructure and controls. Such an approach would also provide redundancy and continuity even if CSP operations were interrupted in-country (despite the above controls) if data and compute embassies were available and coordinated in multiple regions.

THE COPYRIGHT ISSUE

Why hasn't Canada already become a destination of choice for market driven private sector investments in model training infrastructure? There is no indication that the type of GPUs/accelerators installed in country by CSPs are of the kind to power large training runs. As Canada needs to partner with foreign providers for at least some components of its sovereign AI compute strategy, it needs to understand what is preventing them from investing in the country in the first place.

For one, plentiful clean energy supply is in the works, but at this point still in the future (there is currently intense competition for existing energy allocations). Domestic demand for AI compute is also relatively low. Importantly for model developers, Canada's copyright regime does not include an open-ended fair use exemption like in the US (though this isn't a silver bullet), nor is there a clear text and data mining (TDM) exemption as part of the current list of Fair Dealing exemptions. As such, the data scraping associated with model training is on shakier ground in Canada than in most jurisdictions. Combined with statutory fines and the relative ease of filing Canadian class action lawsuits, this makes Canada an unappealing location in which to train models for (non-academic) model developers, whether domestic or foreign. As these are the primary customers of CSPs for AI training infrastructure, there has been little reason to consider investing in the country. Conversely, Japan has a clear TDM exemption and has seen important investments in AI training compute capacity by AWS, Google, Oracle, and Microsoft, despite being significantly more energy constrained.[XIX]

Canada has explored the need for a clear TDM exemption to support AI development in the past, even concluding in the 2019 statutory review of the Copyright Act that the government should pursue such an exemption.[XX] The current political climate will make this difficult, however.[XXI] Without change, this will necessarily limit the attractiveness of Canada for private sector partners to invest in AI training facilities. Should the Government find that it has plentiful dance partners for the SAC deal notwithstanding this issue, then it may leave its resolution to the courts. In the alternative, however, it will need to find some way to address this concern to effectively partner with other countries, CSPs, and model developers. A few potential paths are outlined in Annex 2.

INCREASING THE GOVERNMENT'S OPERATIONAL CAPACITY TO EXECUTE

Executing this strategy will require coordinated action, deep knowledge, and intense groundwork with partners. The new Canadian AI minister should bring added capacity to develop and execute, while the newly announced Major Projects Office could help coordinate financing and the domestic regulatory changes required[XXII] (data and compute embassy legislation, leveraging the Foreign Missions and International Organizations Act, etc.). Still, the team and the Minister are new (and few). The Government should therefore consider adding capacity, including diplomatic capacity, to ensure intense on-the-ground engagement with potential partners. The SAC represents a level of technical, financial, and diplomatic complexity that exceeds the federal government's current institutional and talent base. The Canadian public service today possesses strong analytical and regulatory capacity but lacks sufficient depth in large-scale AI infrastructure, blended finance structuring, and technology diplomacy.

At present, no federal entity is equipped to design or negotiate a public–private venture that integrates high-performance computing, energy infrastructure, and international capital. The government has deep expertise in broadband and energy projects, but not in datacenter dynamics. Likewise, few officials have experience in the structuring of power purchase agreements or long-term compute off-take contracts for artificial intelligence workloads. Canada must also strengthen its ability to navigate export controls and technology-sharing agreements with allies, which increasingly govern access to advanced chips and cloud technologies (and which may be part of CUSMA renegotiations). Finally, this know-how must be coordinated.

A focused institutional response is therefore required. One approach could be to embed a small, mission-driven AI Infrastructure Unit within the Major Projects Office, supported by secondments from the Canada Infrastructure Bank and Export Development Canada. The Major Projects Office could act as the government's control tower for project coordination, regulatory alignment, and interdepartmental approvals. The Canada Infrastructure Bank could bring investment and financing expertise, establishing a dedicated AI Infrastructure Investment Desk to design joint-venture structures, sovereign loan guarantees, and power purchase contracts. Export Development Canada and Global Affairs Canada would provide international structuring and compliance support, particularly for cross-border financing and export control alignment.

Recruitment for this new unit should prioritize Canadians with commercial and technical experience from hyperscalers, semiconductor firms, and infrastructure investment funds, as well as engineers and project financiers currently working abroad who may be attracted to a public mission of national significance. Canada could also draw on short-term secondments from allied governments and multilateral banks or establish temporary advisory roles for senior industry figures under clear conflict-of-interest safeguards. What Canada lacks is not ambition but the specialized human capital to operate at the speed and sophistication of the global AI and semiconductor industries. Building such a capacity within government is the single most important enabler of execution for both the SAC and the CCII.

In parallel, the government should establish an Ambassador-at-Large for AI & Emerging Technology who, if operationally feasible, holds a dotted-line accountability to the AI Minister's office. This model is not without precedent, as Ambassador Fick held a similar role under President Biden. The Minister will retain the strategic leadership of Canada's AI agenda, while the Ambassador-at-Large will serve in an operational capacity, leading direct diplomatic and commercial engagement with key bilateral partners to support the SAC and associated international commitments. This role will focus on the continuous on-the-ground work required to secure cooperation from allied governments and multilateral bodies. While the Minister sets the national priorities, the Ambassador-at-Large will carry out the sustained negotiation of partnership agreements, export-control coordination, and international legal arrangements necessary for success. By embedding this role within the channel of Global Affairs Canada or in close coordination with the Major Projects Office, the government can ensure the Ambassador-at-Large has the rank, access and mandate needed to deliver complex infrastructure diplomacy in support of (rather than duplicating) the Minister's responsibilities.

03 ACCELERATE AI DIFFUSION

History shows that being first to invent a technology is not enough; rather, as Jeffrey Ding's study of past industrial revolutions demonstrates[XXIII], those who "win" geopolitical relevance from the advent of general purpose technologies (GPTs) are those who succeed at diffusing innovations throughout their economy. In his words, technological leadership is "more akin to a marathon" than a sprint, favoring the "patient and adaptive" who incorporate new tools widely, over those who focus narrowly on frontier development of technology.

For example, Britain's dominance in the 19th century owed much to its widespread mechanization and creation of engineering institutions, not just the invention of the steam engine. The United States overtook more innovative rivals in the 20th century largely because it better deployed technologies like electricity and computers across its workforce, supported by broad-based engineering education. The lesson for AI today is clear: the key driver of national power and economic development in the AI era will be how well countries diffuse AI throughout their economies, not necessarily who makes the flashiest breakthrough in a research lab[XXIV]. As Ding notes, "the relative success of nations in diffusing AI…over many decades" will determine who leads, more than any single "AI arms race" to a theoretical AGI.

This provides an opportunity for middle powers like Canada. Canada is highly unlikely to be able to compete on frontier AI development. Yet, as Ding illustrates, when Canada's AI strategy is understood as an investment strategy guided by historical returns, the implications are far from devastating. Rather than being locked out of a race it could not win, Canada can dedicate itself to winning the GPT race that is empirically most likely to matter: diffusion.

The Government should continue and deepen its support for AI R&D and advanced skills acquisition. The issue is not that AI R&D receives too much government funding; it's that research and particularly research talent has been the quasi-totality of Canada's AI strategy. Instead, Canada should broaden its AI investments, not to the detriment of R&D capacity in absolute terms, but to invest a greater share of its strategic analysis, policy development, and new funds towards spurring on diffusion of AI across the economy.

It is widely acknowledged that Canada has faced long-standing issues with productivity growth, both in comparison to its past performance and relative to other advanced economies.[XXV] It's also widely accepted that the diffusion of technologies across firms is a key driver of aggregate productivity growth, that Canada lags its peers in this factor[XXVI], and that GPTs have acted as multipliers of technologically-driven productivity growth.[XXVII]

This isn't news. The Canadian government has acknowledged its productivity problem and the role of technology diffusion in numerous reports, task forces, and Royal Commissions[XXVIII], and it has implemented many of their recommendations. It is not for a lack of trying. Still, it's imperative that the government gets back on its horse and tries again, with a specific focus on AI.

The AI stack has evolved considerably since the 2017 Canadian national AI strategy. AI products and services are more commoditized than ever and are easier to use and access. Many of the requirements for diffusion are now present.

This step-change in AI capacity happens at a moment where Canadian productivity headwinds and tariffs make diffusion a near-term imperative. The central AI policy question should therefore shift from "What is Canada's role in frontier AI?" to "How quickly can Canada make AI do useful work across the real economy?", with the understanding that the returns on getting thousands of firms to adopt proven AI are likely to dwarf the returns from chasing frontier breakthroughs.

While the *What We Heard: Consultations on AI Compute*[XXIX] report consistently flagged affordability as a barrier, this would hopefully be addressed by the compute investments described above. What remains to be developed are the organizational and institutional innovations to spur on AI adoption, like the development of mechanical engineering departments (which trained the skilled workforce for widespread mechanization) or the major factory reconfigurations needed to exploit electrification.[XXX]

A 2025 RSM survey of middle-market firms in the U.S. and Canada found that 63% of organizations feel unprepared for implementing generative AI (despite 91% having adopted it in some form), highlighting that adoption is constrained more by institutional readiness than by lack of access to technology.[XXXIV] A 2025 survey of Canadian AI users found that 64% feel their employers have not provided adequate guidance on how to use AI effectively, with over a quarter (27%) strongly disagreeing that they had received sufficient training.[XXXV]

These findings also concord with a recent MIT study that grabbed the headlines, which found that 95% of AI pilots in enterprises see no measurable P&L impact.[XXXVI] Some have seen this as proof of an AI hype bubble and demonstration that the technology is far from economically transformative. Instead, the study can be interpreted to reveal not problems with AI models that weren't "capable enough", but a learning gap, wherein people and organizations simply did not understand how to adapt to the AI tools, pointing to the need for the "factory layout"-type changes Ding demonstrated were needed with past GPTs.

I suggest a few adoption-aiding policy measures below. It is not an exhaustive list, nor is every policy suggestion fully fleshed out. They build on the work of the AI Adoption Initiative, or "AIAI", an international member-led initiative that aims to network policy practitioners in this space and identify policy best practices for AI adoption.

1- Pick priority sectors for intervention

There's simply no getting around having to struggle with the implementation details to get it right. In our case, this means looking at economic sectors and breaking down "AI" into individual use-cases and applications to find the low hanging fruit for policy-driven adoption and productivity growth. Evidence supports the idea that targeting specific sectors can help speed up adoption rates. Singapore, India, and the UAE, countries with AI adoption rates 50 per cent above the G7 average, have each identified priority sectors within their national AI strategies.[XXXII] Identifying priority sectors can focus policy efforts and help guide how other policy levers, such as workforce skilling and talent policies, R&D support, and sectoral regulation can be designed, targeted and coordinated to optimize for diffusion. Targeting specific sectors can also help mobilize different stakeholders, including universities, the public sector, regulators, businesses, and other institutions that play an important role in AI adoption, to work together on specific programs.[XXXIII]

A specific sector's ranking score could be calculated using the equation:

$$R_i = \frac{(1 - A_i) \cdot S_i \cdot C_i}{N}$$

Where:

R_i is the result score for sector *i*.

A_i represents the sector's level of AI adoption, ranging from 0 to 1 (0 represents no adoption, 1 represents full adoption).

S_i represents the availability of AI solutions relevant for that sector, ranging from 0 to 1 (0 represents no availability of solutions, 1 represents full availability).

C_i represents the comparative advantage of the sector (this can be a value from a predefined scale).

To implement this, I propose to first adopt the sector-selection formula introduced by AIAI. This approach evaluates each sector on three axes: its economic weight, its susceptibility to transformation through AI, and its comparative advantage for the country (which can assign value to national resilience as part of its pre-defined scale). The formula then multiplies these scores to generate a clear priority ranking. The intuition is straight-forward: sectors that are both large and currently transformable, and that carry strategic weight, are where policy interventions will generate the greatest return. This formula offers policymakers a structured way to focus Canada's adoption resources on sectors where AI can drive immediate productivity gains while simultaneously advancing national resilience.

2- Reproduce and improve upon Singapore's Industry Digital Plans for AI

Under SMEs Go Digital, Singapore's Industry Digital Plans (IDPs) provide sector-tailored roadmaps that spell out priority use-cases, recommended digital/AI solutions, and training, paired with a CTO-as-a-Service advisory and grant support (often up to 50%) for pre-approved solutions. In 2024–25, IMDA reported thousands of firms adopting AI-enabled and cloud-based solutions through these channels.[XXXIX]

The Mining Sector and AI [XL]

Mining is central to Canada's economic and geopolitical positioning, especially given the importance of critical minerals for clean energy and advanced technologies. Yet despite the sector's significance, adoption of AI remains medium to low. Nationally, only 12.2 percent of Canadian businesses report using AI to produce goods and services. Surveys of Canadian mines further show that while operators are rapidly investing in sensors and a degree of automation, AI-driven analytics remain under-deployed.

This lag is not due to a lack of technology. Canadian companies already provide proven solutions: Rithmik Solutions (Montréal) offers predictive maintenance platforms for mobile mining fleets; Nanoprecise Sci Corp (Edmonton) provides IoT-based machine fault diagnostics; BigGeo (Calgary) delivers geospatial analytics for exploration and environmental monitoring; and Symroc (Calgary) develops real-time tailings and slope stability monitoring systems. These tools are commercially available and have been deployed in the field. But uptake is uneven and far from universal, especially among mid-tier and remote mines.

The barriers are organizational and institutional rather than technical: integration costs, workforce readiness, regulatory frameworks, and a conservative industry culture. This gap between proven supply and limited diffusion is precisely where an Industry AI Adoption Plan for mining could deliver impact. By subsidizing precise adoption pilots, creating shared adoption roadmaps, and supporting workforce training, government policy can accelerate uptake. This would not only improve productivity and safety but also reinforce Canada's claim to global leadership in responsible critical mineral development.

Canada should launch Canadian **Industry AI Adoption Plans (IAI-APs)** in priority sectors chosen through the above-mentioned AIAI formula. For each sector, the government should (i) publish an applied AI roadmap, detailing individual roles and accompanying tasks, the types of AI applications that can assist with each, and the type of data-sets or tooling required for deployment (e.g. remote sensors or other IOT devices) (ii) pre-approve proven AI solutions that fit the tasks in the IAI-APs, and provide grants/co-funding support to offset deployment and training costs (iii) establish sectoral technical advisors that can be loaned to enterprises to assist with technical deployments. The 100M$ AI Assist Program from IRAP in budget 2024 is a step in this direction[XLI] that could be built upon if the program proves to be a success. In essence, this would be akin to establishing a network of AI extension centers that provide advice and technical support to companies trying to implement AI, much as the US' agricultural extension agents helped farms adopt new crops or techniques.

3- Build an "AI Adoption" talent pipeline

The Government will need a comprehensive and detailed talent strategy for AI that goes beyond deep R&D capacity. It should research and identify each talent bottleneck from model production to SME adoption and address them individually.

For instance, at the highest level, the number of individuals who know how to build foundation models is exceedingly low. Canada is recognized globally as a hub of AI education, training a disproportionate number of graduates in computer science and machine learning relative to its population size and peer countries. However, the "craft" of LLM development (data pipeline design, large-scale optimization, evaluation, and systems engineering) is not taught in academic settings. This is not because the methods are secret or esoteric; rather, they are compute-intensive, operationally complex, and tightly coupled to proprietary industrial environments. As a result, these skills are concentrated within a few, mostly American private firms (OpenAI, Anthropic, Google

DeepMind, Cloud Service Providers and select US research labs) creating a bottleneck in applied AI talent. MILA's suggested LLM school, coupled with the CCII investments described above, would graduate a few hundred additional experts per year to address this bottleneck and unlock a new level of model production capacity.

Beyond top-end research talent, diffusion depends on practitioners who can integrate AI into line-of-business workflows. As the AI stack has become commoditized, understanding how to identify use-cases and how to match these to existing solutions should not require a graduate level degree. Canada should establish an "AI Change Management" professional training program, akin to how the US once rapidly expanded mechanical engineering education during industrialization. This would be a broader, less technically specialized version of the solutions architects described above, who would nonetheless be equipped to work with solutions architects when needed. Setting a clear target could galvanize action and funding; for example, training 10,000 AI-change managers in the next 5 years, embedding them in traditional industries to identify AI use cases and manage implementation. This could take the form of new college programs in AI systems, short-term professional certifications for AI project managers, and up-skilling/re-skilling mid-career workers in AI deployment skills. Just as MBA programs created cadres of managers who spread modern business practices, an *AI management curriculum* (covering technical basics, data governance, change management, use-case identification, etc.) could create a generation of professionals capable of weaving AI into the fabric of Canadian organizations. The government should work with the provinces, potentially through an organization like Mitacs, to establish the program.

Governments (federal and provincial) should also work with educational institutions to embed AI skills across curricula. Every business school, healthcare program, and engineering faculty should include AI modules relevant to their field. For example, an agriculture science student should learn about AI-driven precision farming; a finance student about AI in risk modeling. This cross-cutting education ensures that when graduates enter the workforce, they bring AI awareness into traditionally "non-tech" roles (these types of programs have been successfully executed by Mitacs in the past).

Finally, the government should continue to support K-12 AI education by renewing ISED's CanCode program, currently slated to sunset on March 31, 2026. Since its launch in 2017, youth-serving agencies receiving CanCode funding have delivered digital skills education and training to over 9 million young Canadians and 45,000 educators.[XXXVII] In 2025, AI literacy was integrated into the program. This capacity should be leveraged to broaden AI literacy across the country.

4- Adjust tax policy to support AI adaptation and adoption, not just AI creation

The Government should also examine whether changes to tax policy (and specifically the SR&ED credit program) can encourage faster uptake and diffusion of AI. Recent data from the OECD showed that 33 of 38 OECD member countries offered tax relief for companies that invested in R&D in 2023. Unsurprisingly, these R&D tax breaks are effective at encouraging greater investment in more experimental development, whereas investments in AI adoption, including investments to develop more practical AI applications that can make business processes more efficient, aren't eligible for the same relief.

Tweaking tax policy for AI to recognize the smaller scale experimentation and adaptation required to properly deploy an AI solution makes sense. For example, current tax incentives give companies tax credits for purchases of high-end semiconductors for use in training cutting-edge foundation models, while companies that spend money to fine-tune applications for specific industry use cases -i.e. those working to diffuse AI across specific sectors – are not eligible for similar tax breaks. Tweaking tax policy to allow companies to deduct expenses associated with investments in more applied AI could help spur greater adoption and accelerate progress towards unlocking productivity gains.

5-Lead by example in the public sector (referenceable deployments)

The Government should mandate each major ministry to run at least two AI implementations with rigorous evaluation and publish playbooks/templates, and favour Canadian and open-source tools where feasible to create reference deployments the private sector can trust and emulate.

These reference deployments serve three critical functions to support the Industry AI Adoption Plans (IAI-APs). First, they provide proof of concept and credibility. By demonstrating that Canadian-based AI solutions can operate at scale under public scrutiny, the government helps build trust among private firms that these tools are not only viable but reliable. Successful public deployments send a signal that AI systems can be adapted to Canadian regulatory, cultural, and operational contexts, lowering the perceived risk of adoption for businesses.

Second, reference deployments generate playbooks, standards, and best practices. Rigorous evaluation (including cost, performance, and deployment timelines) creates reusable templates that companies and other levels of government can adopt rather than starting from scratch. This reduces duplication, accelerates diffusion, and ensures that lessons learned in one department can be shared across the ecosystem.

Third, government adoption of Canadian AI tools anchors demand and creates market pull. By integrating domestic solutions into public sector operations, the government establishes a stable baseline of demand that allows local suppliers to scale. This, in turn, reduces costs, improves quality, and strengthens the competitiveness of Canadian firms, ensuring that they can export their solutions abroad with proven use cases at home.

The Government of Canada is already taking steps in this direction. For example, on August 19, 2025, Minister Lightbound (Government Transformation, Public Works and Procurement) joined with the Minister of AI and Digital Innovation to sign a memorandum of understanding with Cohere to explore deploying AI across federal public service operations.

Open-source solutions add a further layer of strategic value. Mandating that at least some reference deployments use open-source systems reduces lock-in risks. It also builds public trust by ensuring that solutions are auditable and modifiable. In addition, open-source outputs from public deployments (whether Canadian-specific language models, sectoral datasets, or evaluation frameworks) can be reused by startups, SMEs, and academic researchers, directly feeding into the IAI-AP sectoral plans. Finally, open-source stacks contribute to resilience: in a geopolitical environment where access to proprietary AI tools may be curtailed, government-supported open systems safeguard Canada against potential supply-chain shocks.

6- Leverage Standards, testing, and evaluation to reduce adoption risk

The government should stand up a light-weight evaluation and accreditation service (in partnership with the Standards Council of Canada and industry) to test vendor solutions for security and reliability and in turn produce short assurance notes to lower perceived risks for SMEs and public buyers.

Importantly, this evaluation must be designed not as a long approval bottleneck but as a speed-friendly assurance path: lightweight, transparent, with clear criteria. It should not be a necessary pre-cursor to adoption or deployment.

By embedding this evaluation infrastructure into reference deployments and IAI-APs, the government can build trust, shorten acquisition cycles, and reduce risk premiums that currently discourage firms from adopting new AI tools. Doing so will help SMEs, public agencies, and less digitally mature sectors see reliable examples and feel confident in moving forward.

7- Promote data readiness and access

Adoption can be stalled by lack of quality data or difficulty accessing data to train models. The government should develop initiatives to improve data *liquidity* in key sectors. For example, creating secure data trusts or pools (also known as "data institutions" [XXXIX]) that companies can draw on to develop AI solutions without compromising privacy or competitiveness. These can take different forms, as outlined by the Open Data Institute, from the UK Biobank, a government-funded body that stores genetic data for research purposes, to Social Science One, which facilitates secure access to data for academics. Such initiatives enable AI innovation across both the private sector and public sector by allowing researchers and practitioners to make use of valuable data in a secure and trustworthy way. In healthcare, Taiwan's public sector Health Passbook is an example of a data institution that has helped build trust to enable personal data access.[XIVL]

In addition to institutional models, crowdsourcing can play a catalytic role in building curated, public-interest datasets. Mozilla's *Common Voice* initiative[XLVI], for instance, has created the world's largest open dataset of human voices by engaging volunteers to donate recordings in multiple languages. Similar approaches could be applied in Canada to domains like Indigenous languages, citizen-science contributions to climate monitoring (images of coastal erosion, invasive species reporting, etc.), urban Infrastructure (community mapping of road conditions, public transport accessibility, or noise pollution levels).

Standards for interoperable data (building upon open banking and open health data frameworks) can also enable smaller players to benefit from large datasets. Additionally, support for data cleaning and labeling efforts (perhaps via incentives or public dataset curation) will help organizations use AI effectively.

Canada already holds globally significant scientific assets in genomics, climate, and health data, yet these remain underleveraged in national AI infrastructure planning. A focused investment in domain-specific datasets could deliver sovereign advantages that compound over time, much as early investments in broadband and satellite data did a generation ago.

One immediate opportunity lies in agricultural genomics. Genome Canada and its partners have built a foundation of world-class excellence in decoding the genetic makeup of crops and livestock that underpin Canada's food systems. Canadian scientists have already sequenced and improved fifteen wheat varieties and developed disease-resistant strains that are saving farmers hundreds of millions of dollars each year. In cattle breeding, genomic tools have generated more than $150 million annually in added value for the beef industry. A coordinated national initiative to build comprehensive genomic databases of Canadian crops such as wheat, canola, and pulses, along with livestock genomes and soil microbiome data, would create a sovereign dataset of immense strategic value. These datasets would enable 190,000 Canadian farm operations to deploy AI tools to adapt to climate change, reduce fertilizer and water use by 20 to 30 percent, and increase yields through precision agriculture. With the global agricultural genomics market projected to reach $12 billion by 2032, Canada has a unique opportunity to turn data leadership into exportable AI and agtech solutions, rather than relying on foreign-controlled platforms.

A second area of opportunity is in health and biomedical data. Canada has one of the world's most trusted public health systems and a strong research base in population health genomics, yet health data remain fragmented across provinces and institutions. While the government is certainly aware of this problem (and organizations like Gemini Health are doing tremendous work in this space), developing a federated, privacy-preserving national health data infrastructure could transform Canada's ability to train AI models for early disease detection, drug discovery, and precision medicine and would be of immense value. Such a dataset would ensure that domestic innovators can develop and validate medical AI models without relying on foreign data.

These strategic datasets should be directly integrated with Canada's broader AI compute infrastructure. The AI Sovereign Compute Infrastructure Program supercomputing facility or the proposed SAC could act as the high-performance backbone for research and model training using these sovereign datasets.

8- Leverage Open-source AI as a diffusion flywheel

A robust open-source AI ecosystem will lead to broader and faster adoption of AI across Canada. Over 50% of organizations are already leveraging open-source AI in each layer of the tech stack (data, models, and tools)[XLII]. Embracing open-source AI aligns perfectly with Canada's needs as a country seeking widespread adoption and heightened resilience and optionality: it lowers costs, reduces dependency on foreign tech giants, and allows solutions to be tailored for local requirements (including on-premises deployment for sensitive data).

The Linux Foundation's 2025 report on open-source AI found that a whopping 89% of AI-using companies rely on some form of open-source AI, with SMEs especially reliant on open models to control costs. Without open source, the report estimates that organizations' spending on AI deployments would be 3.5× higher on average[XLIIII]. These statistics underline that open-source AI isn't just a nice-to-have; it's a major asset in improving AI diffusion, if only through costs savings.

From a competitiveness standpoint, open-source AI also ensures Canada's researchers and companies can stay at the cutting edge. State-of-the-art research tends to be published openly, and implementations often appear in open source within weeks. As one example, when generative image models (like DALL-E) were first announced as proprietary, the open community quickly reproduced similar capabilities (Stable Diffusion and others) that were openly available. This "open chase" of frontier advances greatly shortens the gap between the frontier and the rest. Canada should embrace this dynamic.

However, while the open-source AI ecosystem is vibrant, it can be fragile and under-resourced. Many key tools are maintained by small teams or volunteers with limited resources. We recommend that Canada:

- **Provide Compute and Funding for Open Projects:** One of the biggest needs of open-source AI developers is access to computing power to train and test models. The proposed CCII should explicitly prioritize support for organizations working on open-source AI solutions. Additionally, a sizeable portion of the proposed compute investments and SAC model described above should be allocated to open-source model training projects that align with Canada's goals (such as bilingual language models, healthcare AI, climate modeling tools, etc.). Similarly, grants could be offered through NSERC or Innovate Canada for teams maintaining important open-source AI software, potentially in partnership with like-minded foundations already invested and connected in this space.
- **Adopt Open Standards and Models in Government:** When the public sector develops AI solutions or procurement, it should favour open-source software unless there is a compelling reason not to. This also avoids lock-in and ensures longevity of systems (since open-source tools can be maintained and updated by the community even if vendors vanish).
- **Encourage Open Collaboration in Industry:** The government can act as a convener for industry-specific open-source collaborations. In areas like healthcare AI or smart cities, consortia of companies and researchers could jointly develop core AI models under an open license, sharing the costs and while benefitting all. The Government could seed-fund these efforts and provide legal/organizational frameworks. Sector-specific *AI Innovation Networks* that brings together stakeholders to create common resources (data sets, model benchmarks, etc.) under open principles would accelerate diffusion while avoiding duplicate efforts.

- **Ensure Regulations Support Open Source:** As Canada creates AI-related regulations (for example, around AI safety, accountability, or intellectual property), it should be mindful not to inadvertently smother open-source contributors with undue burden. Large corporations can navigate complex compliance regimes, but volunteer-driven open projects cannot. Likewise, Canada should coordinate with allies to resist any international policy that would broadly restrict open publication of AI research. Canada's strategy should explicitly champion the idea that open-source AI is a force for good driving innovation, education, and economic inclusion and thus should be encouraged, not viewed with suspicion..

9- Cultivating Demand and Culture

Finally, an overlooked aspect of diffusion is creating a culture that *pulls* innovation. Canadian business culture has often been labeled as cautious in adopting new tech[XLIV]. Government leaders and industry champions need to consistently message the importance of embracing AI. Highlight success stories of AI-enabled Canadian firms, create awards or recognition for traditional businesses that have transformed via AI, and disseminate metrics on how AI is boosting productivity where it's used. By celebrating and normalizing the idea that "AI is the new electricity" powering every enterprise, more managers will feel it necessary to jump on board. Quotes like Gudie Hutchings' (then Minister of Rural Economic Development) that "*AI isn't science fiction; it helps businesses get more done – like software that scans lumber to improve mill yield*"[L] are down-to-earth examples that resonate with SMEs. The government and its Ministers need to double-down on relatable narratives to demystify AI.

In practical terms, the federal government can partner with business associations (chambers of commerce, manufacturing associations, etc.) to run AI awareness campaigns and workshops across the country. Small grants could be given to these associations to host "AI days" where local tech providers showcase solutions to local firms. The aim is to seed ideas in business owners' minds about what AI could do for them and to connect them with people who can implement it.

04 CONCLUSION: A MIDDLE-POWER STRATEGY FOR AI LEADERSHIP

Canada's AI future will not be defined by a single moonshot project or regulatory move, but by a balanced strategy that plays to the country's strengths and mitigates its constraints. As a middle power, Canada should neither resign itself to total dependence nor chase unattainable supremacy. Instead, our strategy should be to ensure AI's benefits are felt across every sector and region of Canada through a broad adoption and diffusion strategy while securing the sovereign capacity needed to safeguard our interests.

The Strategic Compute Cluster (SCC) can serve as both a hedge against external shocks and as a platform for domestic AI adoption. In parallel, the Canadian Compute and Infrastructure Initiative (CCII) would support and coordinate the growth of Canada's compute ecosystem, strengthen supply-chain resilience, and develop domestic capabilities in high-performance computing, orchestration, and systems integration, while rapidly providing compute capacity to Canada's advanced AI labs.

The objective is not to match the world's largest AI powers in scale, but to convert strategic infrastructure into enduring national advantage. By linking access to compute with domestic adoption, Canada can create a flywheel of AI investment and domestic productivity growth. The SCC, the CCII and the broader industrial policy measures suggested in this paper can thus become the catalyst for a mature, self-sustaining AI ecosystem, one that powers applications in critical sectors, drives private investment, and embeds Canadian capabilities across

the global AI value chain. In doing so, Canada would position itself not merely as a participant in the global AI race, but as a trusted designer and supplier of democratic AI infrastructure for the world.

Open-source AI will be a force multiplier in this effort, and Canada's policy should unabashedly champion open collaboration. The payoff is a more inclusive AI revolution, and one where world-class AI capabilities can be accessed much more broadly. As the Linux Foundation research highlights , open-source AI reduces costs and accelerates innovation, all while strengthening Canada's strategic autonomy.

Ultimately, Canada's goal should be to drastically improve at getting AI out of the lab and into the real economy. If we do that, our businesses will be more productive, our public services more efficient, and our workforce equipped for the future, translating into economic growth and, as Ding demonstrated, global economic and geopolitical relevance.

This strategy acknowledges our limitations but also our unique advantages. Canada has a rich history of innovation, a strong ability to develop and attract talent, a good reputation with allies, and abundant clean energy potential to power the AI age. By marrying sovereignty with openness, Canada can craft an AI path that is distinctly its own. Canada will not out-spend the superpowers. But as a nimble middle power, Canada can out-adapt and out-cooperate, carving out influence by leading a "third way" of cooperative AI sovereignty.

ANNEX 1

Comparison of Sovereign Cloud Commitments by Major CSPs (AWS, Microsoft Azure, Google Cloud)

This table by Pablo Chavez compares how major cloud service providers implement data sovereignty across five key criteria: data residency, encryption and key management, operational control, regulatory compliance, and transparency. It highlights the varying degrees to which each CSP provides regional data control and governance, illustrating that "sovereign" implementations differ substantially in scope and maturity.

Sovereignty Criteria	AWS	Microsoft Azure	Google Cloud
Data Residency	🟡 Full EU residency for data, metadata, and certificates via new AWS EU legal entity structure. All infrastructure located in EU. Launch in Brandenburg, Germany, by end of 2025. Sources: AWS, DCD	✅ EU Data Boundary active since Jan 2024. In-country options like Bleu ensure data stays local. Sources: MSFT, ITPro	✅ Google Cloud Data Boundary and Sovereign Controls are live. Customers can deploy an EU-only boundary with external key management. Sources: Google, Google Docs
Encryption & Key Management	✅ Strong BYOK with CloudHSM & XKS. Dedicated EU-based root Certificate Authority ensures all cryptographic identity and SSL/TLS certificate issuance remains in Europe. Source: AWS	✅ Full sovereign key management with Azure Key Vault, Managed HSM, BYOK, Double Key Encryption. Source: MSFT	✅ External KMS, client-side encryption, and Key Access Justifications ensure customer control. Source: Google
Local Operational Control	✅ Entirely EU-based operations and governance. New AWS EU parent company and subsidiaries with all leadership and operational staff residing in the EU. Independent advisory board with EU citizens. Source: AWS	✅ Majority EU-owned JVs (Bleu, Delos) with EU staffing and governance; Microsoft provides technology stack but holds minority ownership. Source: MSFT	✅ Majority EU-owned JVs (S3NS, T-Systems) with local staffing and management by Thales and Deutsche Telekom. Sources: S3NS, T-Systems
Operational Autonomy	⚠️ Strong technical autonomy: fully isolated EU infrastructure, routing, and EU-based source code access. However, AWS EU remains U.S.-owned. Data portability available. Sources: AWS, DCD	⚠️ Swiss code escrow enables operational continuity even if Microsoft withdraws service. Azure software and infrastructure remain under Microsoft. Microsoft supports structured data export via Azure Arc et al. Sources: MSFT, MSFT	⚠️ Air-gapped offering available through S3NS enhances survivability. GCP remains owned by Google. Google offers strong data export tools and multi-cloud portability via Anthos et al. Sources: The Register, Google
Regulatory Compliance	🟡 Includes ISO 27001, C5, SOC1/2/3; committed to new Sovereign Requirements Framework (SRF) with independent third-party audits. Signed compliance cooperation agreement with BSI. Not yet SecNumCloud certified. Source: AWS	✅ Strong compliance posture. Bleu seeking SecNumCloud; dedicated EU cybersecurity roles. Source: MSFT	🟡 Certifications (C5, ISO 27001); SecNumCloud pending. Strong partner involvement in compliance. Source: DCD
Transparency & Oversight	🟡 Publishes regular transparency reports that include government data access requests. EU-based advisory board provides governance oversight. Oversight structured through the SRF with third-party audits. Source: AWS	✅ Submits to independent audits and publishes detailed transparency reports covering law enforcement and national security requests. EU board oversight is provided through Bleu and Delos joint ventures. Has a public record of challenging extraterritorial data demands. Source: MSFT	🟡 Publishes transparency reports covering government access requests. Governance of S3NS and T-Systems handled by EU-based partners. Supports Key Access Justifications (encryption-layer visibility). Source: Google

Table: Created by Pablo Chavez. ✅ = Fully implemented, according to publicly available information. 🟡 = In progress, partial, or not yet launched. ⚠️ = Implementation includes strong technical controls, but key legal or structural questions remain. You may reuse or adapt the table as long as you credit me and link back to the original. • Created with Datawrapper

ANNEX 2

Data & Compute Embassies

The concept of the data embassy is an emerging innovation in international diplomacy and digital resilience. Building on principles established under the Vienna Convention on Diplomatic Relations, which grants inviolability to embassies, diplomatic personnel, and communications, the idea extends these protections to data. A "data embassy" applies diplomatic immunity to government data housed in secure, commercial data centres located abroad, allowing a state to retain legal and operational control even when the physical infrastructure is outside its borders. Estonia pioneered this approach after the 2007 cyberattacks, hosting government data backups in Luxembourg to ensure continuity of operations in the event of a domestic IT system failure[XLVII].

For Canada and its allies, this concept can evolve beyond data redundancy to encompass operational AI capacity, e.g. compute embassies that ensure continuity of critical digital functions, including model training and data processing, during crises or infrastructure disruptions. In this model, a Canadian-based Strategic adoption cluster (SAC) could host an "embassy-status" facility shared among trusted partners. Such a facility would operate under a special legal regime recognizing specific datasets and compute workloads as if they were under the jurisdiction of the partner country or of a multilateral organization.

Expanding the embassy model in this way could achieve multiple strategic goals. First, it would pool demand from allied nations, increasing the utilization rate and financial viability of sovereign compute infrastructure in Canada. This pooled approach would make large-scale facilities like the SAC more attractive to private investors and cloud service providers, who could rely on consistent workloads tied to intergovernmental agreements. Second, it would allow governments to circumvent domestic legal or jurisdictional constraints in controlled, legitimate ways, such as enabling AI model training on copyrighted data within a protected environment that is jointly governed and compliant with international norms.

A multilateral compute embassy arrangement could also include conditional access for private-sector entities from participating countries. For example, firms developing models recognized as critical to national security, scientific advancement, or open-source innovation could conduct training within the embassy facility. This framework would provide sovereign immunity for data, workloads, and digital operations, shielding them from local legal claims or export-control restrictions.

ANNEX 3

Copyright & AI

In the late 1990s, Canada introduced a levy on blank audio recording media to compensate rights holders for private copying. The Canadian Private Copying Collective (CPCC) collected a small surcharge on blank media and distributed the proceeds to artists and producers. The levy was designed to make private copying lawful within a managed, collective remuneration framework rather than to prohibit it outright. This approach had a clear legal footing in Part VIII of the Copyright Act and was periodically operationalized through tariffs approved by the Copyright Board and published in the Canada Gazette.[XLVIII]

An adapted version of that framework could help reconcile copyright concerns raised by large-scale AI training in the Strategic adoption cluster (SAC). Policy and legal uncertainty persists regarding whether ingesting copyrighted works for training is permissible or requires licensing, and the federal government has been consulting on options. A targeted levy tied to clearly defined large-scale training activity could provide a collective remuneration pathway for rights holders while giving developers and infrastructure operators greater legal certainty to proceed.[LIV]

A practical design would avoid taxing small players and general compute. Instead, the levy could attach to large commercial training runs above a GPU-hour threshold or to long-term commercial training contracts executed within the SAC. Funds would flow to a collective management entity for distribution to rights holders, analogous to CPCC's role with blank media. The goal is not to meter every datum or work but to recognize aggregate use and share value back to creators in a workable manner.

The strengths of this approach are threefold. First, it offers legal predictability for AI training conducted inside the Cluster by pairing permitted use with compensation, which supports the project's "social licence." Second, it creates a stable, usage-linked revenue stream for rights holders that scales with commercial training activity rather than relying on ad hoc grants. Third, it aligns with existing Canadian practice of using collective remuneration to balance innovation with creator compensation.

There are material challenges that argue for a narrow and carefully scoped design. Defining the levy base for AI is more complex than blank media because training occurs on shared infrastructure and often across borders. Over-broad levies risk causing competitiveness impacts or breaking the SAC's financial sustainability. Historical criticisms of blank-media levies also caution against administrative opacity and delays in distribution[LV]. These risks point to a targeted, threshold-based levy applied at a clear point of collection such as contracted GPU-hours for very large models inside the SAC, coupled with transparent reporting and regular review.

Keeping transaction costs low is essential so that administration does not overwhelm compensation. The mechanism should prioritize simplicity, scale, and automation. Exemptions for academic research (already exempt under the Copyright Act), "true" Open Source model development, and small enterprises would reduce compliance overhead. If the levy is narrow, predictable, and applied only to very large commercial training runs conducted within Canadian sovereign facilities, it can legitimize training for rights holders, reduce litigation risk for developers, and protect the financial model of the SAC.

Recognizing open-source models as valid forms of research under the existing list of copyright exemptions would be another approach. Open source helps Canada build sovereign AI capacity that is auditable, safe if the full training set and training protocol are made open source and less dependent on proprietary foreign systems. Open-source AI also builds legitimacy for public investments and copyright liability exemptions by ensuring they translate into public goods. By explicitly tying AI funding and incentives

to open licensing, Canada could build itself into a trusted hub for public AI, just as Estonia branded itself on e-governance. This would attract researchers, NGOs, and democratic allies concerned about the dominance of closed U.S. and Chinese models.

"Open-Source AI" runs into definitional issues, however, which will likely complicate the application of such an exemption. The government could instead explore creating an intermediary organization that can vet the public benefit nature of the output of the training run and approve them for exemption. For instance, Canada could create a small treaty-based Tri-Sovereign Open-Source AI Institute with a few like-minded countries and use the Foreign Missions and International Organizations Act so that official acts of the Institute (e.g., TDM to train open-source models) are immune from Canadian civil suits unless expressly waived.

REFERENCES

I. Alex Riehl, "Carney Says New Major Projects Office Will Help Build a 'Canadian Sovereign Cloud,'" *BetaKit*, September 11, 2025.

II. Anton Leicht - Datacenter Delusions (substack).

III. Some believe that the race for model development is a race to artificial general intelligence, or "AGI", and that this race is zero-sum (that the first country to AGI will have such capacity to be forever capable of dominating other countries). If you subscribe to this view then investment in true frontier, AGI training capacity is a necessity. Considering the lack of scientific evidence for this argument, and the staggering costs associated with the type of training capacity needed, we do not believe Canada should hedge for this risk by pursuing domestic AGI. Canada has invested in R&D capacity towards this possibility through MILA and Yoshua Bengio's work.

IV. Belcak, Peter, Greg Heinrich, Shizhe Diao, Yonggan Fu, Xin Dong, Saurav Muralidharan, Yingyan Celine Lin, and Pavlo Molchanov. *Small Language Models are the Future of Agentic AI*. arXiv preprint arXiv:2506.02153, June 2, 2025. https://arxiv.org/abs/2506.02153.

V. Chazapis, Alexandros, et al. "Running Cloud-native Workloads on HPC with High-Performance Interconnects." *arXiv preprint* arXiv:2409.16919, 2024. https://arxiv.org/pdf/2409.16919

VI. Chris Mellor, "CoreWeave Tops New GPU Cloud Rankings from SemiAnalysis," *Blocks and Files*, April 3, 2025, https://blocksandfiles.com/2025/04/03/clustermax-gpu-cloud-ratings-and-storage/

VII. Tom Krazit, "Why CoreWeave Thinks Software Wins AI Business," *Runtime*, May 29, 2025, https://www.runtime.news/why-coreweave-thinks-software-wins-ai-business/

VIII. Idem.

IX. Reuters, "Canada Proposed $15 bln Incentive to Boost AI Green Data Centre Investment, Globe and Mail Reports," Reuters, December 12, 2024, updated December 12, 2024, accessed [today's date], https://www.reuters.com/technology/artificial-intelligence/canada-proposed-15-bln-incentive-boost-ai-green-data-centre-investment-globe-2024-12-12/. cio.economictimes.indiatimes.com+5reuters.com+5reuters.com+5

X. Brookfield Asset Management, "*Building the Backbone of AI*," August 6, 2025, accessed October 30, 2025, https://www.brookfield.com/news-insights/insights/building-backbone-ai

XI. Ryan Morrison, "China Expands AI Subsidies with 'Computing Power Vouchers' to Boost SME Adoption," *NotebookCheck*, September 20, 2024, https://www.notebookcheck.net/China-expands-AI-subsidies-with-computing-power-vouchers-to-boost-SME-adoption.1104359.0.html.

XII. Sun Yu, "China Offers Computing Power Subsidies to Support Artificial Intelligence Start-Ups," *Financial Times*, September 20, 2024, https://www.ft.com/content/9d67cda3-b157-47a0-98cb-e8e9842b2c90.

XIII. United States, White House. *America's AI Action Plan*. Washington, D.C.: The White House, July 2025. https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf

XIV. Hawkins, Zoe (Jay), Vili Lehdonvirta, and Luyang Wu. 2025. *AI Compute Sovereignty: Infrastructure Control Across Territories, Cloud Providers, and Accelerators*. Working paper, Oxford Internet Institute.

xv. Richardson, Aris, Haley Yi, Michelle Nie, Simon Wisdom, Casey Price, Ruben Weijers, Steven Veld, and Mauricio Baker. 2025. "How Sovereign Is Sovereign Compute? A Review of 775 Non-U.S. Data Centers." *arXiv* (August 2025). https://arxiv.org/abs/2508.00932.

xvi. U.S. Department of Commerce, Bureau of Industry and Security, *Electronic Code of Federal Regulations*, "15 C.F.R. § 734.9 — Foreign-Produced Direct Product Rule," accessed September 1, 2025, https://www.ecfr.gov/current/title-15/subtitle-B/chapter-VII/subchapter-C/part-734/section-734.9.

xvii. Laura Osman, "Documents Reveal Why Ottawa Can't Stop Using Amazon Web Services," *The Logic*, October 31, 2025, https://thelogic.co/news/ottawa-depends-on-amazon-documents-reveal/

xviii. Chavez, Pablo. "European Digital Sovereignty: U.S. Cloud Company …." *Consensus Drift (Substack)*, June 3, 2025. https://consensusdrift.substack.com/p/european-digital-sovereignty-us-cloud-409.

xix. Amazon Web Services. "AWS to Invest ¥2.26 Trillion in Japan by 2027." *Reuters*, December 12, 2023. https://www.reuters.com/technology/aws-invest-226-trillion-yen-japan-by-2027-2023-12-12/. Google Cloud. "Cloud TPU Pricing." Accessed August 15, 2025. https://cloud.google.com/tpu/pricing. Microsoft. "Microsoft to Invest $2.9 Billion to Accelerate Cloud and AI Transformation in Japan." *News Center*, April 10, 2024. https://news.microsoft.com/apac/2024/04/10/microsoft-to-invest-2-9-billion-to-accelerate-cloud-and-ai-transformation-in-japan/. Oracle. "Oracle to Invest More than $8 Billion in Japan over the Next 10 Years." *Oracle Newsroom*, April 17, 2024. https://www.oracle.com/news/announcement/oracle-to-invest-more-than-8-billion-in-japan-over-the-next-10-years-2024-04-17/.

xx. Canada. House of Commons. Standing Committee on Industry, Science and Technology. *Statutory Review of the Copyright Act*. 42nd Parl., 1st sess. June 2019, Recommendation 23, 60. https://www.ourcommons.ca/DocumentViewer/en/42-1/INDU/report-16/page-60.

xxi. Nicola Heath, "Australian authors challenge Productivity Commission's proposed copyright law exemption for AI," *ABC Arts*, ABC News, August 13, 2025.

xxii. https://www.pm.gc.ca/en/news/news-releases/2025/08/29/prime-minister-carney-launches-new-major-projects-office-fast-track-nation-building-projects

xxiii. Jeffrey Ding, *Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition* (Princeton, NJ: Princeton University Press, 2024).

xxiv. Robert Bellafiore, review of *Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition*, by Jeffrey Ding, The Republic, April 1, 2025.

xxv. Rosell, Carlos; Kaleigh Dowsett; and Nelson Paterson. "A Critical Juncture: Assessing Canada's Productivity Performance and Future Prospects." I*nternational Productivity Monitor*, no. 61 (Fall 2023): 61-92

xxvi. See Réka Juhász, Mara P. Squicciarini, and Nico Voigtländer, "Technology Adoption and Productivity Growth: Evidence from Industrialization in France," *Journal of Political Economy* (forthcoming), quoted in "Technology adoption and productivity growth: Evidence from industrialisation in France," VoxEU, August 13, 2020 for a historical example. See Nick Alexopoulos and Jon Cohen. "Canadian Productivity Growth, Secular Stagnation, and Technological Change." *International Productivity Monitor*, no. 35 (Fall 2018): 3–24. Ottawa: Centre for the Study of Living Standards. https://www.csls.ca/ipm/35/Alexopoulos-Cohen.pdf for example in Canada.

xxvii. Bresnahan, Timothy F., and Manuel Trajtenberg. "General Purpose Technologies 'Engines of Growth?'" NBER Working Paper No. 4148, May 1992 (revised November 2022). National Bureau of Economic Research.

xxviii. See for instance Canada. Royal Commission on the Economic Union and Development Prospects for Canada. *Report of the Royal Commission on the Economic Union and Development Prospects for Canada*, vol. 2, pt. III, ch. 9 "Industrial Policy." Ottawa: Minister of Supply and Services Canada, 1985.

xxix. Innovation, Science and Economic Development Canada. *What We Heard: Consultations on AI Compute*. Government of Canada, March 26, 2025. https://ised-isde.canada.ca/site/ised/en/what-we-heard-report-consultations-ai-compute.

xxx. RSM US LLP and Big Village, *RSM Middle Market AI Survey 2025: U.S. and Canada*, RSM Insights, June 2025, https://rsmus.com/insights/services/digital-transformation/rsm-middle-market-ai-survey-2025.html.

xxxi. TD Bank Group, "Majority of Canadian Workers Are Optimistic about AI but Aren't Ready," news release, August 14, 2025, https://td.mediaroom.com/2025-08-14-Majority-of-Canadian-workers-are-optimistic-about-AI-but-arent-ready.

xxxii. Nathan Furr and Andrew Shipilov, "Beware the AI Experimentation Trap," *Harvard Business Review*, August 29, 2025, https://hbr.org/2025/08/beware-the-ai-experimentation-trap.

xxxiii. IBM Corporation, "*2023 Global AI Adoption Index Report*," IBM Multivu, 2023, accessed September 16, 2025, https://www.multivu.com/players/English/9240059-ibm-2023-global-ai-adoption-index-report/.

xxxiv. Marc E. Ouimette, Ed Teather, and Kevin Allison, *AI, Everywhere, All At Once: A New Policy Agenda for AI Success Through Faster Adoption* (New York: Adopt AI, November 2024), https://adopt-ai.org/wp-content/uploads/2025/04/AIAI_0325_AI-Everywhere-All-At-Once_Updated-Final.pdf.

xxxv. Infocomm Media Development Authority, "Industry Digital Plans (IDPs) | SMEs Go Digital," Singapore Government, accessed August 29, 2025, https://www.imda.gov.sg/how-we-can-help/smes-go-digital/industry-digital-plans.

xxxvi. Statistics Canada, "Canadian Survey on Business Conditions, Q2 2025"; University of Waterloo, "Emerging Technologies in the Mining Industry: Adoption Survey 2023."; *EY Canada, "Top 10 Business Risks and Opportunities for Mining and Metals in 2024"; CIM, "Technology Adoption in Mining: Organizational and Workforce Challenges*."

xxxvii. National Research Council Canada. "*NRC IRAP: Support for SMEs Innovating with Artificial Intelligence*." Government of Canada, accessed September 16, 2025. https://nrc.canada.ca/en/support-technology-innovation/nrc-irap-support-smes-innovating-artificial-intelligence.

xxxviii. Government of Canada. "*Government of Canada Launches Fourth Phase of CanCode Program, Targeting 1.6 Million Training Opportunities for Youth and Teachers*." News release, Innovation, Science and Economic Development Canada, August 7, 2024. Accessed July 23, 2025. https://www.canada.ca/en/innovation-science-economic-development/news/2024/08/government-of-canada-launches-fourth-phase-of-cancode-program-targeting.html.

xxxix. Open Data Institute. "*What Are Data Institutions and Why Are They Important?*" ODI Explainers. London: Open Data Institute, 2020. Accessed September 16, 2025. https://theodi.org/insights/explainers/what-are-data-institutions-and-why-are-they-important/

xl. CEIMIA, "The Role of Governments as a Provider of Data for Artificial Intelligence," CEIMIA Project, accessed September 16, 2025, https://ceimia.org/en/projet/the-role-of-governments-as-a-provider-of-data-for-artificial-intelligence/.

xli. Mozilla Foundation, "*Common Voice: Help Make Voice Recognition Open to Everyone*," Mozilla, accessed September 16, 2025, https://commonvoice.mozilla.org/

xlii. Padmalav Sahoo, *Open Source Technology in the Age of AI* (uploaded document, Scribd, accessed October 30, 2025), https://www.scribd.com/document/871067201/Open-Source-Technology-in-the-Age-of-AI#:~:text=use%20of%20open%20source%20AI,Over%20of%20organizations%20report

xliii. Linux Foundation, *Economic Impacts of Open-Source AI*, research report (San Francisco: Linux Foundation, 2024), https://www.linuxfoundation.org/research/economic-impacts-of-open-source-ai?hsLang=en

xliv. RBC Thought Leadership, "*Bridging the Imagination Gap: How Canadian companies can become global leaders in AI adoption*," The Growth Project, Royal Bank of Canada, June 24, 2025, https://www.rbc.com/en/thought-leadership/the-growth-project/bridging-the-imagination-gap-how-canadian-companies-can-become-global-leaders-in-ai-adoption/

xlv. Government of Canada. "*Federal Government Launches Programs to Help Small-and Medium-sized Enterprises Adopt and Adapt Artificial Intelligence Solutions*," Innovation, Science and Economic Development Canada, October 24, 2024. Accessed September 16, 2025. https://www.canada.ca/en/innovation-science-economic-development/news/2024/10/federal-government-launches-programs-to-help-small-and-medium-sized-enterprises-adopt-and-adapt-artificial-intelligence-solutions.html.

xlvi. *Supra*,

xlvii. Organisation for Economic Co-operation and Development, Observatory of Public Sector Innovation, "Establishing the First Data Embassy in the World," last modified February 7, 2017, OECD OP SI, accessed October 30, 2025, https://oecd-opsi.org/innovations/establishing-the-first-data-embassy-in-the-world/

xlviii. See Copyright Board of Canada. "CPCC – Tariffs for Private Copying, 2018–2019." Certified decision. Accessed October 30, 2025. https://decisions.cb-cda.gc.ca/cb-cda/decisions/en/item/366771/index.do. See also "Tariffs for Private Copying, 2019." Canada Gazette, Part I, Vol. 153, No. 50 (December 14, 2019). https://gazette.gc.ca/rp-pr/p1/2019/2019-12-14/html/sup2-eng.html. decisions.cb-cda.gc.ca+1 and

xlix. Canadian Private Copying Collective (CPCC). "Frequently Asked Questions," and "About the Private Copying Levy." Accessed October 30, 2025. https://www.cpcc.ca/en/frequently-asked-questions and https://www.cpcc.ca/en/wp-content/uploads/2016/10/CPCC_About-The-Private-Copying-Levy_Web_ENG-Oct-2016.pdf

l. See ISED. *Consultation on Copyright in the Age of Generative AI: What We Heard Report*, February 11, 2025. https://ised-isde.canada.ca/.../consultation-copyright-age-generative-artificial-intelligence-what-we-heard-report. Innovation, Science and Economic Dev And Fasken. "Canada's Consultation on Copyright and Generative AI: A Look at the Submissions Received," July 22, 2024. https://www.fasken.com/en/knowledge/2024/07/canadas-consultation-on-copyright-and-generative-ai-a-look-at-the-submissions-received

li. Wired Magazine. "Canadians Burned by Blank-CD Levy," January 2003. https://www.wired.com/2003/01/canadians-burned-by-blank-cd-levy