Building the Stack

From Cloud Dependency to Sovereign Control

Prepared as a strategic roadmap for policymakers, investors, and institutions committed to Canada's sovereign digital future.

By Roy Chartier, Founder & CTO, Qvelo Corporation

Table of Contents

Section *Page*

Section *Page*

Executive Summary

In an era where artificial intelligence is rapidly shaping the foundations of global power, Canada stands at a crossroads. As technological dependency deepens, the sovereignty of data, compute infrastructure, and algorithmic governance becomes a critical matter of national interest. This whitepaper presents an urgent, evidence-backed case for the establishment of a comprehensive *Sovereign AI and Data Infrastructure Strategy* for Canada.

The global acceleration of AI capabilities, coupled with increasing authoritarianism, economic protectionism, and infrastructure weaponization, creates structural vulnerabilities for middle-power democracies like Canada. At present, Canadian public institutions, researchers, and industries rely heavily on foreign cloud platforms, foundation models, and black-box AI services. This dependency is not only a technical liability, it represents a loss of control over the digital public square, national economic competitiveness, and core democratic decision-making processes.

Canada possesses critical advantages: clean energy, trusted public institutions, AI research leadership, and a strong legal tradition. What is missing is a unifying policy, funding, and execution framework that can tie together regional assets, public interest mandates, and industrial policy.

This whitepaper provides a path forward, structured across eight comprehensive sections. Key recommendations include the creation of a *National Sovereign Compute Backbone (NSCB),* a *Sovereign AI Lab Network (NSAIL),* a *Canadian Digital Sovereignty Act (CDSA)*, and the launch of *AI Sovereignty Bonds* and *public-interest model dividends*. It emphasizes investments in public datasets, bilingual open-weight models, and education pipelines that support diverse and inclusive talent development across Canada.

Sovereign AI is not just a matter of national pride—it is a strategic imperative. Countries that fail to control their data, models, and infrastructure will become passive consumers in a digital economy designed by others. Canada has the opportunity—and the obligation—to choose a different path.

— Roy Chartier, Founder & CTO, Qvelo, June 2025

Investor Takeaways

1. *Strategic Thesis:* The next global infrastructure wave will be digital and model based. Sovereign AI infrastructure presents a once-in-a-generation opportunity to build trusted, scalable, and resilient platforms for the 21st century.
2. *Massive Underserved TAM:* Public sector, mid-market, francophone, Indigenous, and nonprofit use cases remain unserved by hyperscalers. Sovereign AI solutions targeting these verticals offer first-mover advantages.
3. *Durable Policy Tailwinds:* AI governance, localization laws, and public sector transparency mandates are proliferating globally. Canadian policy is aligning with sovereign-first digital strategies.
4. *Long-Term Yield Profiles:* AI Sovereignty Bonds and sovereign model dividend schemes offer ESG-compliant, low-volatility, innovation-tied returns.
5. *IP Anchoring:* Open-weight model licensing and Canadian IP trust structures allow for equity retention, monetization, and re-investment into national innovation.
6. *Blended Finance Ready:* The structure allows co-investment by pensions, government agencies, foundations, and infrastructure funds.
7. *Dual Use Markets:* Civilian and defense applications of sovereign AI systems (e.g., wildfire detection, cyber defense, border infrastructure) can be monetized at scale.
8. *Public-Private Catalyzation:* Opportunities to co-develop inference APIs, developer tools, safety layers, and compliance products exist in every Canadian sector.

Introduction:
The Tipping Point for Canadian Digital Sovereignty

Canada is entering a decisive new era—one defined not only by geopolitical uncertainty and economic volatility, but by a global race for control over the most strategic resources of the 21st century: *data, compute, and algorithmic power*. Once seen as a neutral commons, the global internet is fracturing under the strain of digital authoritarianism, platform consolidation, and intensified technological competition. In this new landscape, nations that fail to assert their digital sovereignty risk becoming subordinate actors in someone else's empire—mere tenants on platforms they do not own, governed by laws they did not write.

The stakes for Canada are unusually high. Over the past decade, our public institutions, researchers, and innovators have become deeply dependent on U.S.-based cloud and AI infrastructure. For years, this dependency was rationalized as efficient, affordable, and proximate to global innovation. But that logic has begun to unravel. The 2024 re-election of a radically nationalist U.S. administration—accompanied by executive overreach, tariff shocks, politicized tech regulation, and the reactivation of extraterritorial surveillance powers—has shattered any lingering assumptions about the neutrality or reliability of foreign platforms. The convenience of integration has become a strategic vulnerability.

Canada can no longer outsource sovereignty. Executive orders emanating from Washington now have the power to disable systems, reroute data, or censor platforms used by Canadian citizens, governments, and industries. The CLOUD Act, once theoretical in its implications, is now actively shaping how foreign services treat allied data. Inaction would leave our public sector exposed, our small businesses priced out, and our civic institutions—schools, hospitals, courts—dependent on entities whose interests are neither Canadian nor democratic.

Yet this moment also offers Canada a rare opportunity. Unlike many countries caught between superpowers, Canada brings together an exceptional combination of democratic credibility, rule-of-law institutions, clean energy capacity, and world-class AI research. These advantages position Canada to become a global leader in e*thical, resilient, and sovereign digital infrastructure*—without needing to rely on authoritarian partnerships or geopolitical patronage.

This whitepaper argues that digital sovereignty is no longer optional—it is essential to national resilience, economic independence, and democratic continuity. It lays out the case for building sovereign AI and data infrastructure rooted in Canadian values and public interest principles. It highlights the structural gaps we must address, from talent pipelines to procurement frameworks. And it offers a detailed blueprint for how Canada can move from dependency to leadership through coordinated policy, public investment, and strategic alliances.

The consequences of failing to act will be felt most by those least able to adapt: small businesses, Indigenous communities, nonprofits, public educators, and healthcare workers—groups already facing exclusion from access to advanced AI tools and high-performance computing. Without public infrastructure, we risk locking these institutions into black-box systems governed by profit, not accountability.

We are standing at a digital crossroads. The question is not whether change is coming—but whether Canada will shape that change, or be shaped by it. Moving toward sovereign AI will demand vision, coordination, and capital. But the alternative—permanent digital dependency—is a risk we can no longer afford.

The window is narrow. The need is now. Canada must move from digital dependency to *digital leadership*—or risk becoming a *renter in its own future*.

.

1.0 Strategic Drivers: Why Sovereign Infrastructure Now?

The rationale for Canadian sovereign AI and data infrastructure is not speculative—it is grounded in the accelerating unraveling of geopolitical stability, technological neutrality, and international norms that have underpinned the digital economy for decades. These include the collapse of trust in U.S. platforms, the weaponization of compute and data by hostile and allied powers alike, and the global supply chain crisis surrounding AI hardware. These drivers that make sovereign infrastructure not only necessary but strategically unavoidable.

1.1 Collapse of Trust in U.S. Platforms

Canada's critical infrastructure—spanning health, energy, defense, transportation, and public services—is increasingly digitized, and increasingly dependent on hyperscale cloud platforms operated by U.S.-based multinationals. While these services offer world-class scalability and efficiency, they remain subject to the extraterritorial jurisdiction of U.S. law and the political whims of the executive branch.

The return of a disruptive U.S. administration, marked by arbitrary executive orders and the consolidation of power within national security agencies, has made clear that these platforms cannot be assumed to operate neutrally or indefinitely in the interests of foreign clients. The CLOUD Act, reinterpreted aggressively, allows U.S. intelligence to demand access to foreign-held data if it passes through U.S.-controlled systems. The balance of power has shifted: Canadian public institutions now operate with borrowed infrastructure that can be revoked, surveilled, or manipulated without recourse.

Moreover, recent high-profile cases of platform politicization—such as algorithmic content throttling, unilateral service denial, and selective enforcement of terms of service—illustrate the strategic risks of dependency. Should tensions escalate over trade, intelligence, or sovereignty, these companies may be compelled to enforce national policies, even when they contradict international norms or Canadian law.

1.2 Weaponization of Data and Compute

Data and compute are not just technological assets—they are instruments of power. Nations that control model training infrastructure, inference APIs, and data lakes now wield outsized influence over global norms, knowledge dissemination, and even democratic discourse.

Canada's integration into NORAD and the Five Eyes security alliance places it squarely in the strategic crosshairs of adversarial regimes. Hostile actors are already conducting disinformation campaigns, cyberattacks, and influence operations targeting Canadian institutions. But a more subtle risk lies in AI dependency: when national services rely on foreign AI models and compute platforms, their functionality can be degraded, misaligned, or covertly manipulated.

In 2023, several democratic countries experienced disruptions to legal, healthcare, and government workflows due to upstream model deprecations by U.S.-based firms. Even when well-intentioned, these changes often reflect foreign commercial or legal priorities. Sovereignty in the age of AI means control—not just of data, but of the logic systems that interpret it.

1.3 The AI Arms Race and the Supply Chain Crisis

The global shortage of advanced GPUs and AI accelerators—amplified by geopolitical restrictions on chip exports, natural disasters affecting semiconductor fabrication, and hoarding by hyperscalers—has elevated compute to the status of strategic commodity.

Canada does not possess advanced semiconductor fabrication capabilities, and it is unlikely to develop such capacity at scale in the short term. However, Canada does have abundant clean power, available land, cold climates ideal for cooling, and the institutional expertise to deploy high-density compute campuses.

What is missing is a national strategy: a coordinated investment model to establish sovereign compute zones, a pooled procurement framework to ensure equitable access across provinces and institutions, and a trusted operating system layer for secure orchestration.

Without this, Canadian research, innovation, and national resilience will increasingly depend on foreign actors who may not share our priorities or respect our laws.

2.0 Institutional Strengths: Canada's Competitive Advantage

Canada is not starting from zero. While significant gaps remain in our digital sovereignty architecture, the country enjoys several enduring advantages that uniquely position it to lead in the development of sovereign data infrastructure and trustworthy artificial intelligence. These strengths are not theoretical — they are tangible, underutilized assets waiting to be activated through policy alignment, capital allocation, and institutional coordination.

This section explores six foundational Canadian advantages: world-class AI research institutions; abundant clean energy suitable for compute; democratic governance and legal trust; linguistic and cultural diversity; a strong public education and healthcare base; and existing trade, intelligence, and regulatory partnerships that amplify Canada's voice on the global stage.

2.1 Academic and Research Leadership in AI

Canada's global reputation in AI is not recent, nor is it overstated. Our nation played a foundational role in the rebirth of deep learning through the pioneering work of researchers like Geoffrey Hinton, Yoshua Bengio, and Richard Sutton. The government's early investments in institutes such as Mila (Montreal Institute for Learning Algorithms), the Vector Institute (Toronto), and Amii (Alberta Machine Intelligence Institute) catalyzed a vibrant ecosystem of AI scholarship and startup formation.

These three centers of excellence anchor Canada's sovereign AI potential. Mila has become globally recognized for its leadership in responsible AI and the social impacts of machine learning. Vector excels in both applied research and commercial engagement. Amii has developed a unique niche around reinforcement learning and industrial AI. Together, they represent a deeply rooted intellectual capacity that few countries of Canada's size possess.

What is needed now is a strategic evolution — a shift from isolated excellence to coordinated sovereignty. Canada can transform these institutions into nodes of a federated national AI infrastructure. By pooling GPU access, sharing data pipelines, collaborating on model safety, and unifying procurement standards, these institutes can anchor a secure and democratic alternative to proprietary foreign platforms.

2.2 Abundant, Renewable, and Clean Energy for Compute

Artificial intelligence — particularly large model training — is immensely energy-intensive. The carbon cost of training a single foundation model can exceed that of flying a commercial jet across the Atlantic several hundred times. In this context, countries with abundant, clean, and affordable energy have a foundational advantage.

Canada is one of them. Hydro-powered provinces like Quebec, British Columbia, Newfoundland & Labrador, and Manitoba generate vast surpluses of low-carbon electricity. These regions not only have the capacity to host high-performance compute clusters — they can do so with a sustainability profile that dramatically undercuts competitors.

In addition to environmental advantages, this clean energy provides hard economic leverage. Training a 65-billion parameter LLM in Montreal or Winnipeg can be 40–60% cheaper in electricity terms than equivalent training in California, Ireland, or South Korea.

Moreover, colder ambient temperatures and pre-existing industrial zoning make many Canadian municipalities ideal for sustainable data center development. When paired with district heating systems, AI training facilities could even contribute to local energy efficiency.

A sovereign Canadian AI stack built atop renewable energy would not only be ethical — it would be globally competitive.

2.3 Trusted Governance, Legal Integrity, and International Credibility

Canada’s regulatory environment, rule of law, and policy transparency are among the most respected globally. In contrast to authoritarian regimes that suppress dissent and algorithmically police populations, Canada is seen as a jurisdiction where rights, accountability, and institutional integrity matter.

This perception translates into real strategic advantage. Countries seeking ethical AI partners, jurisdictions for data residency, or reliable model inference providers are already wary of deploying within China or the increasingly unstable U.S. market. Canada has an opening to brand itself as a neutral, democratic platform for sovereign data operations.

This extends into critical trust layers:

- *Data protection regimes:* While imperfect, PIPEDA and forthcoming legislation under Bill C-27 offer a framework for human-centric data control.
- *AI transparency initiatives:* Canada is active in OECD and G7 forums on responsible AI and algorithmic fairness.
- *Public procurement integrity:* Unlike many nations, Canada has functioning disclosure systems, independent oversight, and media accountability.

In a world of digital disinformation and trust collapse, these governance features are not simply reputational — they are monetizable assets. Canada can host, certify, and export trustworthy AI infrastructure the way Switzerland exports financial stability.

2.4 Bilingualism, Multiculturalism, and Dataset Diversity

AI models reflect the data they're trained on. In countries with homogeneous language, culture, or worldview, large models risk being monocultural, biased, or unrepresentative. This introduces real harms: legal advice models that can’t handle Indigenous law, healthcare bots that fail on immigrant dialects, or education assistants that ignore francophone curricula.

Canada, by contrast, is inherently plural. It is officially bilingual. Its urban centers are multicultural. Its provinces administer education, justice, and social services in different ways. This heterogeneity provides a goldmine for building inclusive, representative models.

No other mid-sized country is better positioned to train multilingual LLMs that respect linguistic nuance, cultural safety, and local idioms. Moreover, Canada’s Indigenous nations, with proper consent and data sovereignty agreements, offer the possibility of training AI models that revitalize endangered languages and respect ancestral knowledge systems.

A sovereign AI stack built in Canada can serve not only Canadians but also most people who are excluded by Anglophone-first, Eurocentric model pipelines.

2.5 Public Infrastructure and Safety-Oriented Institutions

While Canada often lags in private-sector innovation scaling, it maintains strong public institutions in healthcare, education, social services, and science. These sectors generate high-quality, structured, and ethically governed datasets — the kind required to train useful, accountable, and human-centered AI models.

For example:

- *Healthcare:* Provincial EMR systems and CIHI datasets provide rich, longitudinal patient data (if privacy-respecting frameworks are adopted).
- *Education:* Ministries across provinces generate standardized testing, curriculum, and administrative data that could power tutoring and assessment models.
- *Social services:* Structured intake, casework, and policy analytics systems offer training data for decision support tools.

Crucially, these sectors are subject to public oversight. This makes them safer environments for deploying, testing, and refining AI models under tight regulatory conditions — a sandbox for responsible innovation.

In contrast to the "move fast and break things" ethos of Silicon Valley, Canada can become the global leader in AI that is safety-tested, democratically governed, and civic-embedded.

2.6 Global Alignment Opportunities

Finally, Canada is exceptionally well-placed to coordinate with other like-minded nations on digital sovereignty efforts. Through the Five Eyes alliance, OECD, G7, and bilateral agreements with the EU, Japan, South Korea, and India, Canada can act as both a convenor and contributor to international AI norms.

Examples include:

- *Trustworthy AI consortiums* with European data jurisdictions
- *Federated LLM training initiatives* with India and France
- *Joint GPU infrastructure buildouts* with Nordic countries
- *Common open-weight and licensing standards* for publicly funded models

By taking the lead in embedding sovereignty, safety, and transparency into infrastructure layers, Canada can define the next generation of internet architecture — one that resists authoritarianism and empowers democratic innovation.

3.0 Vulnerabilities: What Canada Must Address

While Canada's strengths in AI research, clean energy, and public institutions are formidable, they are counterbalanced by critical structural vulnerabilities that must be addressed if the country is to chart a viable course toward digital autonomy. These vulnerabilities span hardware and software supply chains, funding models, platform dependency, data residency, legal and procurement frameworks, and the absence of centralized coordination.

Without urgent redress, these gaps will not only stymie progress—they will deepen Canada's dependency on foreign infrastructure, making sovereignty increasingly difficult to reclaim over time. This section identifies and analyzes Canada's seven most acute vulnerabilities in the context of sovereign AI and data infrastructure.

3.1 Absence of Domestic Semiconductor Fabrication

The foundation of modern AI infrastructure is hardware: specifically, advanced semiconductors for training and inference acceleration. Canada does not have a single operational foundry capable of manufacturing high-end chips at the 5nm or 3nm nodes required for state-of-the-art GPUs and AI ASICs.

This creates a cascading risk. Canadian compute capacity is wholly dependent on the U.S. and East Asia (primarily Taiwan, South Korea, and increasingly, China) for chip design, fabrication, and assembly. Supply shocks—from geopolitical conflict, trade restrictions, or natural disasters—would instantly compromise Canada's ability to scale or maintain sovereign compute resources.

Moreover, Canada's limited participation in global chip alliances (e.g., CHIPS Act, EU Chips Act) reduces its access to advanced nodes and restricts its bargaining power in multi-nation procurement efforts.

While full semiconductor independence may be infeasible, a sovereign AI strategy must at minimum:

- Participate in multi-lateral chip sharing and joint ventures
- Encourage Canadian-led design of specialized accelerators
- Invest in back-end packaging, testing, and integration capacity
- Stockpile critical components and diversify sources of GPU acquisition

3.2 Over-Reliance on U.S. Cloud and SaaS Platforms

Canada's public sector—including universities, hospitals, municipalities, and federal ministries—remains overwhelmingly dependent on U.S.-based cloud providers (AWS, Microsoft Azure, Google Cloud) and SaaS platforms (Salesforce, Workday, Oracle, etc.) for both compute and enterprise software.

This introduces multi-layered risk:

- *Jurisdictional Exposure:* Data processed on U.S. platforms, even when stored in Canadian regions, is subject to U.S. law (notably the CLOUD Act).
- *Service Continuity Risk:* Political volatility in the U.S. could result in disruptions, embargoes, or data seizures—especially for sensitive sectors like defense, immigration, and health.
- *Innovation Lock-In:* Canadian institutions are often contractually bound to proprietary APIs, model architectures, and support ecosystems, limiting their ability to migrate to open alternatives.

This platform dependency is particularly dangerous in AI. If foundational models are served exclusively via U.S. APIs, Canada has no ability to audit, modify, fine-tune, or align them to Canadian legal or ethical norms.

The only remedy is structural: development and national support for sovereign cloud services, open-source AI model infrastructure, and domestic alternatives for critical public workloads.

3.3 Fragmented Digital Infrastructure Governance

Canada lacks a unified digital sovereignty authority or coordinating body responsible for aligning investments, setting standards, and auditing compliance across provinces and sectors.

Instead, digital infrastructure decisions are dispersed across:

- Provincial Ministries of health, education, and energy
- Crown corporations
- Academic institutions
- Federal departments (PSPC, ISED, Treasury Board Secretariat, etc.).

This results in inconsistent procurement standards, duplicative infrastructure spending, fragmented cybersecurity policies, and lost economies of scale. For instance, three major provinces may independently negotiate AI cloud service agreements with different vendors, none of which include provisions for model explainability or data export controls.

A sovereign strategy requires a central coordinating function with cross-jurisdictional authority—akin to how Shared Services Canada consolidated federal IT infrastructure, or how Infrastructure Canada guides long-term capital projects.

This entity should set minimum AI safety standards, manage sovereign cloud accreditation, and coordinate national GPU allocation for research and public services.

3.4 Insufficient GPU Access for Canadian Researchers

Despite leading the world in AI theory, many Canadian labs cannot train or run modern foundation models because they lack access to advanced GPUs like the NVIDIA H200 or AMD MI300X.

Why? Several reasons:

- *Pricing:* Market costs for such GPUs, especially in cloud environments, are prohibitively high for academic budgets.
- *Scarcity:* Most advanced GPUs are locked up by hyperscalers for internal LLM development or sold via exclusive long-term contracts.
- *Procurement Lag:* Public sector entities often cannot acquire AI hardware fast enough due to red tape and funding cycle constraints.

The result is that world-class Canadian researchers are forced to use outdated hardware, apply for limited compute time on U.S. systems, or outsource work to foreign platforms.

This asymmetry is eroding Canada's AI leadership. Without a national GPU reserve or shared sovereign training cluster, the country will soon lag not only in production capability but in the foundational science that supports ethical AI.

3.5 Legal Grey Zones in Data Ownership and Residency

Canadian data law—while progressive in many respects—remains outpaced by the complexities of AI training, inference, and cross-border data flows. Key gaps include:

- No clear framework for the ownership of data-derived models (e.g., who owns a model trained on Indigenous medical data?)
- Limited enforcement capacity for violations of data residency requirements
- Incomplete protections for inferred data, synthetic data, or downstream model behavior.

Additionally, few legal mechanisms exist for requiring transparency in how models are trained, how datasets are curated, or how outputs are validated.

This undermines public trust, deters private-sector innovation, and invites misuse of personal or public datasets by foreign platforms operating under non-Canadian norms.

To resolve this, Canada needs a comprehensive AI-augmented update to privacy law, along with data fiduciary provisions and legally binding audit rights over third-party models operating in regulated sectors.

3.6 Procurement Policies that Favour Foreign Vendors

Canada's own procurement frameworks, particularly in public sector technology, unintentionally reinforce dependency on foreign platforms. Key issues include:

- *Lowest-bid bias:* Contracts are often awarded to vendors with the lowest up-front cost, regardless of long-term sovereignty or risk
- *Non-Sovereign RFP criteria:* Many calls for tender are written in a way that do not disqualify U.S. based providers (e.g., citing FedRAMP alignment instead of explicitly requiring PBMM)
- *Lack of digital value multipliers:* Canadian bidders offering open source or domestically hosted options are not credited for strategic benefits like data control or local employment

This procurement myopia locks Canadian institutions into high-risk, foreign-controlled ecosystems.

Sovereign AI requires a new procurement doctrine: one that emphasizes long-term resilience, interoperability, and national value retention alongside cost. Precedents exist—such as Canada's CRTC decisions in telecom infrastructure—and should be mirrored in AI and cloud policy.

3.7 Talent Leakage and Brain Drain

Finally, Canada's investment in top-tier AI talent is not translating into sovereign capacity because of ongoing brain drain to foreign platforms.

Graduates of Canadian AI PhDs are routinely hired by U.S. and Chinese firms who offer 2–3x salary, equity participation, and first access to foundational model architectures. In many cases, Canadian-trained researchers are working on closed models that cannot be deployed back into Canada due to export control laws or contractual IP restrictions.

This brain drain is not just economic—it is strategic. Without retention, Canada loses institutional memory, research continuity, and the ability to train the next generation of sovereign technologists.

Policy responses must include:

- AI-specific immigration pathways to attract global talent
- Mission-oriented grants and sovereign model fellowships
- Guaranteed infrastructure access for academics building open models
- Employee stock ownership and revenue sharing incentives for talent working in sovereign-aligned initiatives

4.0 Infrastructure Priorities: Building Sovereign AI Foundations

To secure its digital autonomy, Canada must move beyond aspiration and into implementation. This section lays out the infrastructure priorities necessary to support a national sovereign AI strategy—spanning compute capacity, data stewardship, model development, physical security, network architecture, and trusted orchestration.

These components are not abstract. Without them, Canadian researchers, public institutions, and industrial AI efforts will remain tethered to foreign cloud infrastructure and opaque black-box models. Each priority detailed below represents a concrete, fundable initiative that—if executed in concert—would establish the backbone of a Sovereign AI future.

4.1 National Sovereign Compute Backbone (NSCB)

The first priority is the development of a geographically distributed but federated network of high-performance AI clusters across Canadian provinces, collectively known as the *National Sovereign Compute Backbone (NSCB)*.

Each NSCB site should be designed to:

- Host 1,000–10,000+ GPUs (H200/MI300 class);
- Operate on provincially sourced clean energy (hydro, nuclear, wind, geothermal)
- Be interconnected via high-bandwidth (400+ Gbps) national fiber and research backbones
- Include secure air-gapped zones for government-classified workloads and national security research
- Support hybrid access models (e.g., pay-per-use, reserved capacity for public institutions, and "donor credits" for open research).

The NSCB should be governed through a *federated operational model,* where each site is locally managed (e.g., by a university, provincial crown corporation, or regional innovation hub) but coordinated nationally through shared API standards, GPU reservation protocols, and cybersecurity frameworks.

This approach prevents a single point of failure, encourages local economic development, and allows rapid scaling through public-private co-investment.

4.2 National Model Repository and Trust Registry

A Sovereign AI strategy must include infrastructure for model provenance, integrity, and deployment tracking. Canada should establish a *National Model Repository and Trust Registry (NMTR)*—a secure platform where all publicly funded models and datasets are catalogued, versioned, and assessed.

Key capabilities should include:

- Secure storage of open-weight models (e.g., LLaMA, Falcon, northern-trained variants)
- Watermarking and lineage tracking to document dataset usage, training parameters, and updates
- Cryptographically signed validation of model authenticity and compliance with Canadian legal standards
- Model cards with explainability metadata, bias audits, and alignment benchmarks
- Federated lookup for courts, regulators, journalists, and citizens to verify the use of AI in decisions affecting them.

This NMTR would operate akin to a “digital standards body,” establishing a baseline for safe, auditable model use in public-facing systems. It would also support commercial players seeking to demonstrate ethical model stewardship.

4.3 Sovereign Data Lakes and Sector-Specific Data Meshes

Data is foundational to AI—but Canada's most valuable datasets are fragmented, siloed, and often underused. The country must invest in a *layered data infrastructure strategy* that includes:

- *Sovereign Data Lakes:* Nation-scale repositories for high-value, low-risk public datasets (e.g., census, climate, agriculture, scientific papers). These should be curated, permissioned, and queryable via secure APIs.
- *Sectoral Data Meshes:* Federated data exchanges for regulated domains like healthcare, education, finance, and justice. These must respect privacy and jurisdictional controls, but enable anonymized AI training, simulation, and model refinement.

Examples:

- A HealthMesh enabling AI training on distributed patient data without centralizing PII.
- An EduMesh allowing tutoring AIs to be trained across curricula in all provinces.
- A LegalMesh enabling law model pretraining on bilingual court decisions, legislation, and administrative rulings.

Canada's investment in public services has generated decades of rich structured data. With the right infrastructure and governance, these datasets can power a sovereign AI ecosystem that reflects Canadian realities—not foreign abstractions.

4.4 Trusted Orchestration Layer: Canadian AI Operating System (CAIOS)

Above the hardware and data layers lies the need for a *trusted orchestration layer:* an open, auditable, secure AI platform stack that coordinates access to compute, data, and models across Canadian institutions.

We propose the creation of *CAIOS: the Canadian AI Operating System,* a modular orchestration platform built on open-source foundations (e.g., Kubernetes, MLflow, Ray, Slurm, Airflow) with Canadian-specific extensions for:

- Multi-tenant GPU job scheduling
- Role-based access control (RBAC) for model deployment
- AI ethics and audit plug-ins
- Model tuning and evaluation dashboards
- Compliance layers for Canadian privacy and accessibility laws.

CAIOS would allow researchers, public institutions, and startups to safely run AI workloads without being tied to hyperscale cloud APIs or opaque model-serving interfaces. CAIOS should be developed as a public-private partnership, ideally through a combination of university labs, provincial innovation hubs, and certified national integrators.

4.5 Secure Data Centers and Edge AI Zones

To physically host this infrastructure, Canada will need to designate and develop a tiered network of *Sovereign Data Centers (SDCs)* and *Edge AI Zones (EAZs)*.

- *SDCs* should meet strict criteria:
    - Located in geopolitically secure regions
    - Powered by clean, redundant energy sources
    - Hardened against cyberattack, EMP, and fire
    - Certified for multi-tenant use with zero-trust access architecture
    - Monitored by Canadian security-cleared personnel
- *EAZs* should be distributed at provincial and municipal levels and include smaller-scale GPU and storage nodes for:
    - Local inference (e.g., police bodycam analysis, smart city sensors)
    - Emergency resilience (e.g., air-gapped model deployment in disaster scenarios)
    - Regional model fine-tuning and customization

By blending national-scale SDCs with regional EAZs, Canada can reduce latency, preserve bandwidth, support diverse linguistic contexts, and ensure graceful degradation under network partition or supply chain stress.

4.6 Sovereign Network Interconnects

AI infrastructure requires not just physical hardware and compute—but also *network transport sovereignty.* Canada must ensure the confidentiality, availability, and resilience of its digital arteries.

Key initiatives include:

- Expansion and protection of *NREN (Canada's national research and education network),* ensuring all sovereign AI nodes can move petabyte-scale datasets with minimal latency.
- Secure optical backbones interlinking SDCs and EAZs, preferably with *Canadian-owned fiber,* rather than leased or routed via U.S. exchange points.
- Dedicated *interprovincial sovereign fiber corridors,* particularly in underserved Northern, Indigenous, and rural regions.

Additionally, backbone capacity should be ring-structured to withstand regional outages and enable deterministic routing for sensitive model and data flows.

4.7 Integration with Government and Civic Use Cases

Finally, sovereign AI infrastructure must be integrated into *real-world applications* that directly benefit Canadian citizens. Examples:

- Legal aid chatbots trained on Canadian law
- Open-source curriculum tutors aligned with provincial education standards
- Healthcare triage models built from anonymized provincial EMR data
- Public weather and climate models hosted on sovereign compute
- Civic participation models for multilingual engagement with local governance

These projects both justify infrastructure spend and build public trust by showing that sovereign AI is not just a strategic asset—it's a civic utility.

5.0 Regulatory Roadmap: Legal and Policy Foundations for Sovereign AI

Sovereign AI is not just a technological or infrastructure problem, it is a legal and institutional one. Without the appropriate regulatory scaffolding, Canada risks building powerful AI systems that operate in legal grey zones, reproduce systemic harms, or fail to deliver public benefit. Worse, without strong protections and controls, Canada may find itself enforcing foreign standards against its own citizens while its own public institutions remain beholden to closed foreign AI platforms.

This section outlines the critical policy reforms, legislative frameworks, and enforcement mechanisms required to enable a secure, rights-preserving, and economically productive sovereign AI environment in Canada.

5.1 National AI Rights and Responsibilities Framework

At the foundation of Canada's sovereign AI regime must be a *National AI Rights and Responsibilities Charter*, similar in spirit to the Canadian Charter of Rights and Freedoms but focused on the emergent digital context.

This framework should codify a set of enforceable rights for individuals and duties for AI deployers, including:

- The *right to opt out* of AI decision-making in high-stakes contexts (e.g., employment, healthcare, justice)
- The *right to meaningful explanation* of decisions made or informed by AI
- The *right to know* when an AI is being used
- The *right to contest* AI outputs and demand human review
- The *right to protection from synthetic manipulation* (e.g., deepfakes, fake news bots, algorithmic targeting)
- The *responsibility of deployers* to conduct risk assessments and fairness audits
- The *responsibility to maintain human oversight* and prevent algorithmic monoculture

This Charter would be enforceable across all AI systems deployed by federally regulated entities and strongly recommended (via funding incentives) for adoption at the provincial and municipal levels.

It would also provide the legal predicate for Canada's participation in international AI rights harmonization efforts—e.g., under the OECD, G7, and UNESCO AI principles.

5.2 Binding Model Transparency and Registration Requirements

Canada should mandate that all *foundation models used in critical sectors* (e.g., government, healthcare, finance, education, law enforcement) be:

- Registered with a federal AI registry
- Accompanied by detailed documentation including training data provenance, known biases, testing results, and intended use-cases
- Benchmarked for compliance with Canadian accessibility, bilingualism, and human rights standards
- Auditable by an independent AI Safety Board (analogous to the Auditor General for algorithms)

This would apply both to public sector deployments and to private-sector vendors serving government agencies or regulated industries.

The regulatory framework should further require:

- *Pre-deployment impact assessments* for high-risk AI applications
- *Continuous logging* of inputs/outputs for redress mechanisms
- *Mandatory human-in-the-loop safeguards* for critical functions (e.g., arrest decisions, border screening, welfare eligibility).

This transparency regime should be built into procurement templates, federal RFPs, and digital services standards.

5.3 Open Weight Requirements for Public-Funded AI Models

To prevent dependency on foreign black-box models, any AI model *developed with Canadian public funds* should be subject to *open weight licensing unless* a strong national security or privacy exemption applies.

This principle would:

- Ensure reusability and retrainability of Canadian AI assets
- Reduce redundant procurement and vendor lock-in
- Support a vibrant open-source AI community within Canada
- Enable third-party validation, error correction, and civic oversight

Exceptions should be narrowly tailored and adjudicated through *an AI Use Exemption Board*—a cross-functional body that evaluates proposed closed deployments for public risk.

The public should always have access to models that are built with their tax dollars and affect their lives.

5.4 Data Sovereignty and Digital Fiduciary Reform

Canada's privacy legislation—PIPEDA—is no longer fit for purpose in the age of AI. A sovereign AI policy must be undergirded by *comprehensive digital fiduciary reform,* which reframes personal data not merely as a commodity, but as a form of delegated agency.

Key provisions should include:

- Explicit data residency requirements for sensitive sectors (e.g., health, child services, national security)
- Strong consent requirements for model training on PI or inferred data
- Prohibitions on synthetic profile creation without transparency and redress
- The establishment of *Data Fiduciaries*—licensed entities who steward data on behalf of citizens with legal duties of care and loyalty
- Enforcement teeth for violations, including algorithmic disgorgement and revocation of licenses for non-compliant vendors

This fiduciary model offers an ethical counterweight to extractive surveillance capitalism and aligns with emerging European trends (e.g., the Data Governance Act, AI Act).

5.5 AI-Specific Antitrust and Procurement Reform

Concentration in the AI sector—particularly among U.S. and Chinese hyperscalers—has reached existential proportions. Canada must proactively update its *antitrust and procurement frameworks* to:

- Prevent AI vendor lock-in for public institutions
- Encourage multi-vendor, interoperable cloud and model deployments
- Penalize abusive terms of service that restrict auditing, customization, or portability
- Prefer open architectures and Canadian-hosted alternatives in public-sector RFPs

Public dollars must not reinforce foreign monopolies. Canada's procurement system must evolve to reflect *digital sovereignty premiums*—e.g., giving weight to:

- Canadian data storage
- Domestic talent retention
- Reinvestment into the Canadian innovation ecosystem
- Commitments to bilingualism and accessibility

The government's own digital services agency (CDS) and Shared Services Canada should pioneer new procurement templates that reflect these values.

5.6 AI Safety Standards and Certification Regimes

Alongside legal protections, Canada must develop *AI safety standards* similar to how it certifies nuclear, aviation, or pharmaceutical technologies.

We recommend the establishment of a *Canadian AI Safety Standards Board,* mandated to:

- Develop risk tiering schemes for AI systems
- Set certification requirements for model training practices, validation, and behavior under adversarial conditions
- Issue guidance for red teaming, model interpretability, and fail-safe modes
- Publish national safety baselines for sector-specific AI systems (e.g., health AI, legal AI, policing AI)

Over time, the Board should maintain a *Canadian AI Safety Index,* publicly rating large models on attributes such as robustness, bias mitigation, hallucination rates, and compliance with national ethical standards.

This ensures that trust is not a marketing claim—it is a certified, inspectable status.

5.7 Cross-Border AI Governance and Treaty Participation

Canada must assert a clear international role in shaping AI governance. It should:

- Join and shape global standard-setting organizations (e.g., ISO/IEC AI committees);
- Lead negotiations for an *AI Non-Proliferation Treaty,* particularly around autonomous weapons and mass disinformation models
- Establish *bilateral sovereign AI partnerships* with the EU, Japan, South Korea, India, and the African Union
- Advocate for *mutual digital sovereignty clauses* in trade agreements to ensure Canada can localize data and models without being sued by foreign firms under investor-state dispute systems

Canada should not only react to AI geopolitics—it should shape it, anchored by values of fairness, transparency, and rights preservation.

5.8 Model Attribution and Liability Reform

Finally, a critical regulatory innovation will be the creation of *AI Attribution and Liability Protocols,* including:

- Mandatory model attribution headers (e.g., “Response generated by GPT-CAN v1.2, Vector Institute, 2026”) for AI-generated content in high-trust domains
- Shared liability frameworks between model developers, deployers, and fine-tuners in case of harm
- Mandatory insurance coverage or indemnification for commercial model deployments
- A public AI Harms Tribunal where citizens can seek redress for misinformation, discrimination, or algorithmic fraud

This would bring AI closer to the accountability standards we demand from pharmaceuticals, automakers, or banks—sectors where product failure results in concrete liability.

6.0 Talent Strategy: Scaling Sovereign AI Expertise in Canada

Any long-term Sovereign AI strategy lives or dies by its ability to cultivate, retain, and deploy talent across all levels of society. Canada's world-leading reputation in AI is the result of deliberate investment in research (e.g., CIFAR, NSERC, the Canada First Research Excellence Fund), a strong immigration system, and early backing of institutions like MILA, Vector Institute, and Amii.

Yet despite these advantages, Canada is hemorrhaging its AI talent to U.S. big tech, Chinese startups, and EU research centers. This loss is not inevitable—it is a symptom of structural gaps in our national talent pipeline, incentive systems, and lack of domestic AI infrastructure.

6.1 AI Education and Training: Scaling Up the Talent Pipeline

To maintain sovereign control over AI systems, Canada must not only graduate more AI specialists—it must do so across disciplines, regions, and career stages. We propose the following actions:

A. Expand AI-Centric Curricula Across the Education System

- *K–12:* Introduce foundational AI concepts (ethics, logic, algorithmic bias, data representation) into digital literacy curriculums, starting as early as grade 6.
- *Post-Secondary:* Fund interdisciplinary AI programs that combine computer science with law, medicine, environmental science, Indigenous studies, and humanities. Embed responsible AI coursework as a graduation requirement in technical fields.
- *Community Colleges:* Develop AI technician programs in regional colleges focused on dataset annotation, model evaluation, and operations—creating "blue-collar" AI jobs aligned with economic inclusion goals.

B. National AI Fellowship Programs

- Launch a Canadian AI Corps offering graduate fellowships for students working on open-source, bilingual, or public-interest AI models.
- Offer sovereign AI training grants to students from underserved communities, Indigenous populations, and rural regions.
- Create a "Sovereign AI Postdoc Bridge" to fund research in priority national areas (e.g., bilingual LLMs, environmental modeling, Indigenous data sovereignty).

C. Credential Recognition

- Fast-track recognition for AI-relevant degrees from trusted international institutions, paired with local credential bridges and ethics courses.

6.2 Retention and Attraction: Reversing the Brain Drain

Canada loses hundreds of highly skilled AI professionals each year to better-funded, stock-compensated roles in U.S. hyperscalers and Chinese startups. This brain drain undermines public investment in education and weakens sovereign capacity.

To reverse this trend:

A. National Sovereign AI Lab Network (NSAIL)

- Establish a distributed national network of Sovereign AI Labs (analogous to Germany’s Fraunhofer Institutes or the U.S. DOE labs) where talent can work on state-aligned, open AI challenges.
- Offer tenure-like stability for researchers to pursue foundational, interdisciplinary, or high-impact public AI work without needing to chase venture capital or corporate buyouts.

B. Competitive Public Compensation Models

- Introduce tiered sovereign AI fellowships with compensation competitive to the private sector.
- Enable revenue-sharing or IP monetization for open-weight models developed in Canadian labs.
- Pilot stock-option-like incentives tied to the economic impact of sovereign AI deployments (e.g., language tutors, health triage, smart agriculture).

C. AI Immigration Acceleration

- Create a dedicated *Global AI Visa Track* with 2-week processing for high-potential researchers, engineers, and builders.
- Offer family relocation support, guaranteed cloud access, and integration into regional sovereign AI hubs upon arrival.

6.3 Workforce Deployment: Matching Talent with National Missions

Sovereign AI talent must not concentrate solely in a few elite institutions or urban tech corridors. Canada's strategy should intentionally distribute AI expertise across domains and geographies.

A. National AI Deployment Challenge Grants

- Launch competitive funding rounds (akin to the DARPA model) for solving national-scale AI problems in:
    - Wildfire prediction
    - Indigenous language preservation
    - Automated legal triage for underserved Canadians
    - Energy grid optimization
    - Climate change modeling

Winners receive direct funding, compute credits, and long-term infrastructure support through the NSCB (National Sovereign Compute Backbone).

B. Provincial Talent Anchor Programs

- Co-fund AI faculty and researchers at provincial universities and colleges under the condition that their work supports regional industries (e.g., fisheries, mining, agriculture, forestry, northern logistics).
- Embed Sovereign AI fellows within provincial public health units, legal aid organizations, and school boards.

C. Public-Sector AI Talent Exchange

- Establish a two-way AI secondment program between federal departments (e.g., CRA, Health Canada, Transport Canada) and AI institutes/startups.
- These roles should count toward tenure, public pensions, and career advancement tracks, encouraging cross-pollination between public and private ecosystems.

6.4 Civic Literacy and Societal Readiness

Sovereign AI is not only about PhDs, it's about public trust, social alignment, and ethical guardrails. To avoid backlash, misuse, or disenfranchisement, Canada must equip its citizens with tools to understand and navigate the AI systems affecting their lives.

A. Civic AI Curriculum

- Co-develop non-technical AI literacy programs with libraries, unions, Indigenous communities, and newcomer organizations.
- Topics: how algorithms work, identifying disinformation, understanding your data rights, how to contest automated decisions.

B. Media Literacy and AI Verification Infrastructure

- Launch national education campaigns to help citizens distinguish between real and AI-generated content, especially ahead of elections.
- Pair this with a *Trusted AI Content Verification Platform*, accessible via browser extensions and social platforms.

C. AI Ethics Councils and Community Forums

- Create community-level AI advisory boards (analogous to local ethics review boards), empowered to audit municipal deployments, host public consultations, and escalate misuse.
- Fund cross-cultural forums on AI's role in society—particularly among Indigenous, francophone, LGBTQ+, and disabled populations.

6.5 Talent Infrastructure: Tools, Resources, and Enablers

Talent development is not just about people—it's about the environment they operate in. To enable productivity, Canada must ensure the following are universally accessible to researchers and builders:

A. Open Access Compute

- Ensure students and startups have access to sovereign GPU clusters without paywalls or restrictive licensing.
- Allocate a percentage of NSCB capacity to "Compute Grants" for low-resource but high-impact AI projects.

B. Data Commons

- Expand national open datasets (weather, mobility, government publications, climate models) and establish clear licensing for research reuse.
- Encourage provincial data pools to align with national metadata and API standards.

C. Sovereign DevTool Chains

- Fund development of Canadian-hosted open-source developer tools for model evaluation, dataset curation, training pipelines, and ethical scoring.
- Ensure tools are bilingual and accessible to visually and cognitively diverse users.

6.6 Measuring Impact and Closing the Feedback Loop

Finally, talent strategy must be measured. Canada should create a Sovereign AI Talent Observatory to track:

- AI degree output and employment by province
- Gender, racial, and socioeconomic diversity in AI fields
- Model attribution and retention (e.g., how many Canadian-trained researchers contribute to foreign black-box systems vs open Canadian models);
- Regional economic uplift from AI deployments
- AI job market volatility, gaps, and emerging needs.

This Observatory can also issue an *Annual Sovereign AI Workforce Report Card*, which will inform federal/provincial education policy, immigration strategy, and research funding allocations.

7.0 Capital Formation and Public Investment Strategy

Building sovereign AI capacity is not just a matter of principle, it is a capital-intensive industrial transformation on par with 20th-century railway expansion, national electrification, or the postwar science complex. Sovereign AI infrastructure will not arise from research grants alone. It requires a durable, multi-decade investment regime: coordinated across public finance, private capital, philanthropic support, and returns from national-scale AI deployments.

This section outlines how Canada can marshal capital to build sovereign AI infrastructure, empower domestic companies, reduce long-term AI dependency, and generate sovereign wealth from its digital assets. It draws from lessons in national broadband deployment, cleantech stimulus, sovereign wealth fund models, and DARPA-style innovation financing.

7.1 Public Capital Vehicles: Anchoring the Buildout

A. National Sovereign AI Infrastructure Fund (NSAIF)
Canada should establish a *public investment vehicle*—similar to Canada Infrastructure Bank (CIB)—to directly finance sovereign AI assets:

- High-performance GPU clusters (via co-investment with provinces and utilities)
- Sovereign data centers and compute zones
- AI safety and ethics centers
- Secure fiber interconnects between institutional partners
- Model and dataset registries
- Bilingual open-weight model training initiatives

Initial capitalization should be at least *$3–5 billion*, drawing from:

- Federal economic development budgets (ISED, NRCan, DND, Health)
- Strategic Innovation Fund (SIF) reallocations
- Green bonds and ESG-aligned public infrastructure loans
- Potential reinvestment of carbon tax revenue into green AI infrastructure

NSAIF investments should include *blended finance structures*—e.g., loan guarantees, matching grants, convertible infrastructure notes—that crowd in private capital while retaining public strategic control.

B. Matching Provincial Co-Investment Agreements
NSAIF should operate in partnership with provinces via matching funds. For example:

- Québec funds sovereign data lakes aligned with francophone and health mandates
- BC co-funds West Coast sovereign AI compute tied to LNG, forestry, and indigenous climate modeling
- Ontario backs model explainability centers aligned with legal and health AI

This ensures alignment with regional economic priorities and creates political durability.

7.2 Mobilizing Private Capital: Sovereign AI as an Investable Asset Class

To scale impact, sovereign AI infrastructure must attract *private institutional capital*—including pension funds, insurance pools, and sovereign-aligned venture capital.

A. AI Sovereignty Bonds
Canada should issue *Sovereign AI Bonds*—green-tech-style government-backed instruments with yields tied to:

- Carbon savings from inference performed on green sovereign compute (vs U.S. cloud)
- Canadian IP revenue generated via open-weight models licensed abroad
- In-kind contributions to public model training (e.g., companies donating GPUs or datasets)

These bonds could be targeted toward:

- PSP Investments
- Ontario Teachers' Pension Plan
- Canadian banks' impact funds
- Indigenous trusts
- Family offices and ESG-aligned foundations

B. Sovereign AI Investment Tax Credits (SAITC)
Introduce a refundable tax credit (e.g., 25–35%) for:

- Capital expenditures on Canadian AI model training
- Upgrades to sovereign data centers and GPU racks
- Contributions to sovereign data commons
- Employment of Canadian AI safety professionals, ethicists, and accessibility engineers

This would build on the success of the Scientific Research and Experimental Development (SR&ED) credit, with a specific sovereign digital lens.

C. Public-Private Training Infrastructure Partnerships
Encourage banks, telcos, pharma, mining, and agriculture sectors to:

- Co-develop vertical-specific sovereign AI applications (e.g., climate risk models, AI mineral surveyors)
- Train models on private-public data mesh layers
- Host sovereign model inference as part of regulated services

Such partnerships should include rights for the Crown to reuse or replicate derivative models for the public good, via standardized IP arrangements.

7.3 Intellectual Property (IP) Retention and Monetization

Canada invests heavily in AI research, but much of its value is exported through foreign IPOs, data monetization by U.S. platforms, and academic spinouts acquired by hyperscalers.

To retain value domestically:

A. Create a Sovereign AI IP Trust
Establish a Crown-controlled trust that:

- Holds IP for all public-funded AI models, datasets, and training pipelines
- Issues open-source and dual-licensing frameworks
- Collects usage metrics, manages attribution, and oversees commercial sublicensing
- Reinvests licensing revenue into future sovereign AI training grants

B. Canada First AI Licensing Clause
Any company receiving over $5M in Canadian public funding for AI R&D must:

- Grant a non-exclusive license to the Sovereign AI IP Trust
- Ensure model weights and training pipelines can be run on sovereign Canadian infrastructure
- Prioritize commercialization pathways that include Canada-based inference hosting and model customization

C. IP Buyback and Repatriation Fund
Create an “IP Repatriation Vehicle” to acquire, adapt, or fork high-impact AI models currently hosted offshore but co-developed by Canadians—particularly where these models serve the public interest.

7.4 AI Revenue Sharing and Model Dividend Models

Sovereign AI models should generate value—not only for companies but also for the Canadian public.

A. Public Model Dividend
For sovereign foundation models trained on Canadian data and compute, implement:

- Per-inference micropayment schemes from commercial licensees
- Annual dividend disbursements to the Sovereign AI Fund from licensing revenue
- Tiered public access based on use-case (e.g., free for education, nominal for SMBs, higher for commercial SaaS)

B. Community Compute Credit Programs
Model dividend revenue can be returned as:

- Compute credits for civic tech, education, and Indigenous programs
- Subsidies for public-sector and non-profit AI delivery and experimentation (e.g., AI-enabled legal aid)
- Free access to validated "AI toolkits" for Canadian nonprofits and social services

This creates a virtuous cycle: public investment → model training → public benefit → reinvestment in the ecosystem.

7.5 Integration with Sovereign Wealth and ESG Portfolios

Canada's sovereign wealth and public pension entities hold over $1.5 trillion in capital. A fraction allocated to sovereign AI would yield:

- Long-term social dividends
- AI job creation in Canada
- Future-proofing against foreign digital dependence
- ESG-aligned returns

Example Investments:

- CPPIB acquires a 30% stake in a sovereign data center operator
- AIMCo funds a sovereign AI agricultural modeling startup focused on prairie provinces
- BCI co-invests in edge compute infrastructure for public health AI deployments

Sovereign AI is no longer a speculative thesis—it is a competitive investment thesis with compound national returns.

7.6 Governance and Risk Management

All capital formation must be paired with high-trust governance and risk oversight.

A. Independent Sovereign AI Investment Board (SAIIB)
A national governing body with representatives from:

- Treasury Board
- Canadian AI research institutes
- First Nations innovation councils
- ESG funds
- Provincial science ministries

Responsibilities:

- Approve sovereign infrastructure buildouts
- Monitor returns and public benefit impact
- Sanction or unwind investments in non-compliant projects
- Publish annual Sovereign AI Investment Reports

B. Risk Reserves and Contingency Planning
Allocate 5–10% of all sovereign AI capital budgets to:

- Resilience measures (cyber, EMP, supply chain)
- Open-weight model red-teaming and retraining
- Lawsuit defense funds for defending model licensing and open-source access

7.7 Measuring Impact and Return on Sovereign AI Capital

To attract capital and sustain political momentum, Canada must measure:

- Jobs created in AI infrastructure, services, and safety
- Public-sector cost reductions via sovereign AI deployments
- Avoided cost from foreign SaaS licensing and data exit
- IP monetization and reinvestment loops
- Greenhouse gas reductions from inference performed on low-carbon sovereign compute
- Sectoral AI adoption levels (e.g., law, education, logistics, energy)

Metrics should be publicly available via a *Sovereign AI Dashboard*, updated quarterly.

8.0 A National Action Plan for Sovereign AI and Data Infrastructure

Considering the strategic, economic, and geopolitical imperatives outlined in this whitepaper, Canada must move decisively to establish and defend its digital sovereignty. Sovereign AI is not a niche technology project, it is a foundational pillar of 21st-century statehood, akin to energy independence, monetary sovereignty, or national defense. The following policy recommendations are designed to provide a coherent, actionable roadmap for federal and provincial governments, academic institutions, and industry partners.

Each recommendation aligns with one or more of the preceding chapters, synthesizing technical, legal, economic, and governance strategies into a unified national agenda.

8.1 Enact the Canadian Digital Sovereignty Act (CDSA)

Objective: Legislate a comprehensive framework that protects Canadian control over critical digital assets—compute, data, infrastructure, and AI systems.

Key Provisions:

- Declare sovereign compute, data residency, and AI decision-making as national interests.
- Mandate public-sector compliance with Canadian data localization and compute sovereignty rules.
- Establish the Canadian Sovereign Digital Infrastructure Registry, listing accredited cloud/data/model platforms approved for federal and provincial procurement.
- Define extraterritorial legal protections for Canadian data stored or processed abroad.
- Fund enforcement mechanisms through ISED and the Treasury Board Secretariat.

8.2 Establish the National Sovereign Compute Backbone (NSCB)

Objective: Create a pan-Canadian network of high-performance sovereign compute infrastructure to support public-sector AI training and inference.

Key Actions:

- Deploy regional GPU compute centers in partnership with universities and clean energy utilities.
- Connect compute sites with dedicated fiber to enable federated model training.
- Provide baseline compute access to researchers, non-profits, startups, and Indigenous communities.
- Integrate access controls, logging, and AI-specific monitoring capabilities.
- Create a National Compute Allocation Board to oversee resource distribution by mission priority.

8.3 Launch a National Sovereign AI Lab Network (NSAIL)

Objective: Build a distributed research and development network focused on Canadian-prioritized AI models, datasets, and safety protocols.

Structure:

- Hubs in key provinces, co-funded by the federal government, research councils, and regional governments.
- Specializations by region: health AI (Toronto), francophone language AI (Montreal), environmental modeling (BC), Indigenous data (Prairies/North), legal AI (Ottawa), logistics/agriculture AI (Atlantic).
- Public access to open-weight models, annotated datasets, and training tools.
- Strong commercialization pathways with Crown IP retention or open licensing.

8.4 Enforce Open Weight and Public Auditability Requirements for Public-Funded AI Models

Objective: Ensure transparency, retrainability, and verifiability of AI systems funded by Canadian taxpayers.

Policies:

- All models developed with >50% public funds must publish weights, training data documentation, and fine-tuning pipelines unless exempted.
- Third-party audit rights granted to AI safety centers or ethics boards.
- “Reproducibility by design” mandates for federal agencies deploying AI systems.
- Mandatory disclosures for any foreign code libraries or pretrained weights embedded in production systems.

8.5 Establish the Canadian AI Safety and Ethics Board

Objective: Create a national agency empowered to regulate and certify AI models, platforms, and practices.

Mandate:

- Develop sector-specific risk tiers (e.g., medical AI vs. social media bots).
- Issue safety certifications for foundation and application models.
- Operate a public AI model rating index based on robustness, bias, and hallucination rates.
- Mandate AI incident reporting and manage a sovereign AI “recall” mechanism for dangerous models.
- Serve as liaison to international bodies (OECD, ISO, G7 AI) to shape global norms.

8.6 Issue Sovereign AI Bonds and Establish the Sovereign AI Investment Fund

Objective: Mobilize long-term capital to fund sovereign digital infrastructure and public-interest model development.

Financial Instruments:

- Sovereign AI Bonds (10–20 year maturity), linked to ESG and innovation indices.
- Blended finance structures for GPU center deployment, sovereign cloud, and data common initiatives.
- Anchor LP investments into Canadian AI startups aligned with sovereign compute and model hosting.
- Revenue-sharing models tied to inference volume, licensing, or carbon savings from local inference.

Governance:

- Fund administered independently with dual reporting lines to Finance Canada and ISED.
- Provincial co-funding mechanisms tied to regional AI priorities.

8.7 Create the National Digital Talent Strategy for Sovereign AI

Objective: Build a resilient, inclusive, and mission-aligned AI workforce.

Components:

- AI curriculum integration from K–12 through to graduate and professional programs.
- Public sector AI fellowships, including placements in health, legal, education, and defense ministries.
- Fast-tracked AI visas and immigration support packages.
- Regional AI career bootcamps and technician certification programs in underserved regions.
- Public interest AI corps for Indigenous, climate, and francophone initiatives.

8.8 Institute a Digital Data Commons Framework

Objective: Provide open, governed access to high-quality, annotated Canadian datasets for sovereign AI development.

Deliverables:

- Launch the Canadian Data Commons Platform with tiered access based on sensitivity and credentialing.
- Develop data licensing standards aligned with Indigenous data sovereignty principles (e.g., OCAP®).
- Fund data annotation jobs across Canada in partnership with colleges and remote communities.
- Align metadata and ontology standards with provincial and municipal open data portals.

8.9 Reform Procurement and Establish Digital Sovereignty Scoring in RFPs

Objective: Incentivize Canadian organizations to select sovereign AI platforms and models.

Reforms:

- Add “Digital Sovereignty Index” scores to all federal RFPs for digital services, models, and cloud platforms.
- Score based on: data residency, model transparency, local IP ownership, Canadian language support, accessibility, and retrainability.
- Pre-qualify vendors offering open-weight models, Canadian-hosted inference, and Canadian retrain pipelines.
- Penalize submissions that rely exclusively on black-box foreign platforms.

8.10 Coordinate Provincial and Federal Regulatory Alignment

Objective: Prevent fragmentation and inconsistency across Canada's AI governance landscape.

Mechanisms:

- Launch the Sovereign AI Federal-Provincial Working Group.
- Harmonize procurement standards, licensing rules, and auditability requirements.
- Allow provinces to opt into shared compute, IP trust, and public model dividend programs.
- Facilitate interprovincial inference routing and data sharing through trusted exchange networks.

8.11 Build International Alliances for Sovereign AI

Objective: Collaborate with like-minded democracies to shape the norms, APIs, and treaties of the AI era.

Actions:

- Negotiate mutual digital sovereignty treaties with the EU, Japan, India, South Korea, and African Union.
- Establish joint open-weight model research with global labs (e.g., GRAIL, EleutherAI).
- Develop bilateral GPU and AI accelerator sharing agreements with trusted partners.
- Push for a UN or G20 declaration on AI non-proliferation, deep-fake governance, and public auditability standards.

8.12 Launch the Sovereign AI Dashboard and Public Metrics Portal

Objective: Ensure continuous public visibility and accountability over Canada's sovereign AI initiatives.

Metrics Tracked:

- Compute availability and utilization across sectors
- AI talent creation, retention, and wage distribution
- Public inference volume on sovereign vs foreign platforms
- Licensing revenue from open-weight models
- Regional access to compute, data, and training tools
- Environmental impact of sovereign AI infrastructure

This portal should be maintained by the Treasury Board Secretariat in partnership with NSAIL and updated quarterly.

Appendices

Appendix A – Key Metrics to Monitor

- *Sovereign GPU cluster utilization by sector* refers to the measurement and reporting of how national high-performance compute resources are allocated and consumed across key Canadian sectors—such as healthcare, education, defense, research, and climate science—to ensure transparency, efficiency, and alignment with national priorities.
    - Understanding how Canada’s sovereign GPU infrastructure is used across sectors is essential for ensuring equitable access, mission alignment, and strategic return on public investment. *Sovereign GPU cluster utilization by sector* refers to the real-time and historical measurement of compute usage allocated to key domains such as healthcare, education, public safety, climate science, legal services, and commercial innovation. By tagging and analyzing compute workloads through secure orchestration layers, policymakers and infrastructure operators can monitor who is using sovereign compute, for what purpose, and at what scale.
    - This data enables Canada’s *National Compute Allocation Board (NCAB)* to make evidence-based decisions about capacity planning, workload prioritization, and emergency allocations (e.g., wildfire modeling or pandemic response). It also supports reporting requirements for ESG disclosures, academic benchmarking, and public transparency. Over time, sectoral utilization data can inform pricing models, reveal unmet demand in underserved regions or disciplines, and shape targeted infrastructure investments. Critically, this metric ensures that Canada's sovereign AI infrastructure remains a *public good—not a privileged asset—accessible and responsive to national priorities*.

- *Open-weight model licensing revenue and adoption* tracks the economic value and usage rates of publicly released Canadian AI models—measuring how open-weight sovereign models are adopted by industry, academia, and public institutions, and how their licensing contributes to national innovation and reinvestment through public IP frameworks.
    - Open-weight AI models—particularly those trained with public funds or sovereign datasets—represent a powerful mechanism for both maximizing accessibility and generating sustained economic value. Open-weight model licensing revenue and adoption refers to the financial returns and usage penetration achieved through structured, tiered licensing of publicly released Canadian models across academia, industry, and government.

Under this model, core foundation models developed through the Sovereign AI Lab Network or other public channels are made openly accessible for non-commercial or civic use, while commercial deployments are governed through *dual-license frameworks* that return royalties or usage fees to the public domain.

- The *Sovereign AI IP Trust* plays a central role in managing these licenses—tracking where and how models are deployed, attributing lineage, and collecting usage-based revenue. This revenue is reinvested into future model training, safety audits, and compute subsidies for open research, forming a *self-reinforcing innovation cycle*. Adoption metrics also provide insight into sectoral demand, ecosystem maturity, and global competitiveness of Canadian AI assets. By embracing open-weight licensing as a strategic lever, Canada can simultaneously advance transparency, reduce dependency on opaque foreign systems, and generate long-term, mission-aligned economic returns from its sovereign digital infrastructure.

- *Public sector vs commercial inference compute volume* measures the proportion of sovereign AI infrastructure used by government, education, and public institutions compared to private sector entities—providing insight into equity of access, return on public investment, and the strategic distribution of national compute resources.
    - As sovereign AI infrastructure comes online, it is essential to monitor how compute capacity is distributed between public sector missions and commercial use cases. Public sector vs commercial inference compute volume tracks the proportion of GPU time, model deployment, and inference activity allocated to government departments, educational institutions, healthcare systems, and public research, versus private-sector actors using the same sovereign infrastructure. This metric provides insight into whether Canada's AI resources are fulfilling their intended role as a national utility—serving the public good—while also enabling commercial innovation in a balanced, transparent manner.
    - This information is crucial for the National Compute Allocation Board (NCAB), which uses these metrics to assess equity, adjust pricing models, and prevent resource monopolization. Ideally, a baseline allocation—e.g., 40–60%—would be reserved for public-purpose inference, with the remainder monetized through commercial workloads that also contribute financially to system sustainability. Over time, this utilization metric supports informed policymaking, targeted subsidy design, and the tracking of mission-aligned outcomes such as reduced government IT outsourcing, improved accessibility, and public trust in AI systems. It also helps demonstrate fiscal responsibility by showing how public compute delivers public value.

- *Geographic distribution of AI talent and education investment* analyzes how AI-related human capital, training programs, and funding are allocated across Canada's regions—highlighting disparities, informing regional development strategies, and ensuring that sovereign AI capacity is built inclusively across urban, rural, Indigenous, and underserved communities.
    - Ensuring that AI talent and education funding are equitably distributed across Canada is critical to building a truly sovereign and inclusive national AI ecosystem. Geographic distribution of AI talent and education investment refers to the tracking and analysis of how human capital, institutional support, and public funding for AI-related programs are allocated across provinces, territories, and urban-rural divides. Without deliberate intervention, AI expertise tends to cluster around a few major research hubs—leaving northern, Indigenous, rural, and smaller urban communities underrepresented in both opportunity and infrastructure access.
    - This metric informs national policy and workforce development efforts by identifying talent deserts, underfunded institutions, and high-potential regions lacking AI educational programs or job pathways. By aligning data with outcomes—such as AI job creation, research output, and diversity indicators—Canada can target its investments more effectively through programs like the National Digital Talent Strategy, provincial anchor faculty initiatives, and sovereign AI technician training programs in community colleges. Over time, equitable geographic investment builds resilience, fosters inclusive innovation, and ensures that *sovereign AI is not just concentrated—but distributed across the country it is meant to servE*.

- *IP monetization and reinvestment flow* tracks how intellectual property generated through publicly funded AI models and infrastructure—such as open-weight models and sovereign datasets—is licensed, commercialized, and reinvested into Canada's research ecosystem, ensuring long-term value capture and innovation sustainability.
    - A well-structured Sovereign AI ecosystem must not only enable innovation but also capture and reinvest its economic returns. Central to this is the establishment of a *Sovereign AI IP Trust*—a public-interest entity tasked with holding, licensing, and managing intellectual property generated through publicly funded AI models, datasets, and training infrastructure. Rather than allowing high-value Canadian AI assets to flow into proprietary foreign platforms or corporate consolidations, the IP Trust ensures that these outputs remain accessible, reusable, and economically productive within Canada's innovation ecosystem.
    - Monetization occurs through *tiered licensing models,* including open-weight licenses for academic and civic use, dual licensing for commercial applications, and sovereign-retained licensing for critical infrastructure. Revenue from these models—whether through recurring licensing fees,

inference metering, or derivative model royalties—is funneled back into a reinvestment loop. This loop funds future model training, compute subsidies for underserved researchers, safety and alignment audits, and new public-interest model development. Over time, the IP Trust acts not only as a guardian of national AI assets but as a *flywheel for sustained innovation and sovereign capability*, reinforcing Canada's autonomy in the global digital economy.

- *Carbon offset from sovereign vs foreign inference* measures the emissions reduction achieved by running AI workloads on Canadian sovereign infrastructure—powered by low-carbon energy sources—compared to foreign cloud platforms, quantifying the environmental benefits of localized, green AI compute.
    - One of the strategic environmental advantages of sovereign AI infrastructure is its ability to deliver measurable carbon offsets by localizing inference workloads on low-emission Canadian compute. Unlike many global hyperscale cloud providers whose data centers may rely on fossil-intensive energy mixes, Canada's sovereign AI clusters can be powered by hydroelectric, nuclear, and other clean energy sources. By shifting inference—from large language models, vision models, and public-sector AI services—onto domestic infrastructure, each workload execution avoids a significant amount of upstream emissions. This delta can be quantified and reported as a *carbon offset,* contributing to national climate goals and enhancing the ESG profile of AI deployments in healthcare, education, and research. Over time, a standardized methodology for measuring and certifying these offsets could support new financing instruments—such as *green AI bonds*—and embed environmental accountability into public sector AI procurement and infrastructure planning.

Appendix B – Suggested Governance Bodies

- Sovereign AI Investment Board (SAIIB)
    - The *Sovereign AI Investment Board (SAIIB)* would be an independent, multi-stakeholder body responsible for overseeing Canada's public and blended investments in sovereign AI infrastructure, ensuring alignment with national priorities, transparent capital deployment, and long-term public benefit.
- Canadian AI Safety and Ethics Board (CAISEB)
    - The *Canadian AI Safety and Ethics Board (CAISEB)* would be a national regulatory body tasked with certifying AI models, setting safety standards, overseeing ethical compliance, and conducting independent audits to ensure that AI systems deployed in Canada are robust, transparent, and aligned with Canadian legal and societal values.
- National Compute Allocation Board (NCAB)
    - The *National Compute Allocation Board (NCAB)* would be responsible for managing access to Canada's sovereign GPU and AI compute infrastructure, allocating resources based on mission priority, public benefit, and research impact across sectors such as healthcare, education, climate, and defense.
- Sovereign AI Federal-Provincial Working Group (SAIFPWG)
    - The *Sovereign AI Federal-Provincial Working Group (SAIFPWG)* would be a coordinating body that facilitates alignment between federal and provincial governments on sovereign AI policy, infrastructure deployment, regulatory standards, and public-sector adoption strategies, ensuring cohesive national execution while respecting jurisdictional responsibilities.

Appendix C – Additional Reading and Global Benchmarks

- European Commission: AI Act and Gaia-X Infrastructure
    - https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
    - https://gaia-x.eu/
- India: Bhashini Multilingual AI Stack
    - https://bhashini.gov.in/bhashini-at-work
- United States: Executive Order on Safe, Secure, and Trustworthy AI (2023)
    - https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence
- Japan: Trusted Web White Paper
    - https://www.kantei.go.jp/jp/singi/digitalmarket/pdf_e/Trusted_Web_White_Paper_ver3.0_Overview-1.pdf

Appendix D – Model Candidates for Sovereign Training

- Francophone Canadian LLM (bilingual + legal + health tuned)
    - *The Francophone Canadian LLM* would be a bilingual large language model specifically trained on Canadian French and English, with domain-specific tuning in legal, healthcare, and public-sector datasets to ensure linguistic equity, regulatory compliance, and culturally aligned AI services across Canada's official languages.
- Indigenous knowledge graph and translation models
    - *The Indigenous Knowledge Graph and Translation Models* would be AI tools co-developed with Indigenous communities to preserve, map, and translate Indigenous languages, oral histories, and cultural frameworks—supporting language revitalization, digital inclusion, and self-determined data governance in line with OCAP® principles.
- Environmental forecasting ensemble models (fire, flood, crop)
    - *The Environmental Forecasting Ensemble Models* would be sovereign AI systems designed to simulate and predict high-impact climate events—such as wildfires, floods, and agricultural yield variability—by integrating multi-source Canadian geospatial, meteorological, and ecological data to support emergency response, infrastructure planning, and climate resilience.
- Legal and regulatory summarization of AI for courts and agencies
    - *The Legal and Regulatory Summarization AI* would be a sovereign model designed to assist Canadian courts, tribunals, and public agencies by generating concise, bilingual summaries of legislation, case law, regulatory filings, and administrative decisions—enhancing accessibility, reducing workload, and supporting transparent, AI-assisted legal workflows.