8

Canadian Shield Institute

Foundations of Digital Sovereignty
Chapter 8 - June 2026

Sovereign Compute: Options for Canada's Digital Infrastructure

Contents

Contact:

Canadian Shield Institute
150 King St. West,
Toronto, Ontario, CA

canadianshieldinstitute.ca
info@canadianshieldinstitute.ca

Authors:

Vass Bednar
James McLeod
David Corbett
Emily Osborne
Kaylie Tiessen
Matthew da Mota

Photography:

Hermes Rivera
Manny Fortin
Kris Tian
James McLeod

Overview

Throughout the previous chapters of *Foundations of Digital Sovereignty,* we have examined why governance is the most important factor for asserting Canadian sovereignty over the digital realm.

On its own, building data centres on Canadian soil will not do much to bolster our digital sovereignty. Clear regulation on data usage is a necessary precondition to set guardrails for what happens inside the data centres and what digital services ultimately get delivered to Canadians.

Alongside data regulation, we need a better strategic approach to intellectual property, trade agreements, technical standards, procurement, and other systems that allow the government to shape the digital economy and capture value for Canadians.

Without a governance layer, any tax dollars spent on building new data centres on Canadian soil will ultimately leak value to foreign technology giants, and we will still struggle to govern the digital economy in a meaningful way.

However, if we accept the premise that compute capacity and data storage will be essential to the future economy, then it is in the national interest for Canadians to have access to reliable cloud compute and cloud storage.

Some taxpayer-funded infrastructure can provide reliable compute capacity for researchers, public service delivery, and innovative Canadian businesses. Government funded infrastructure can strategically fortify governance of cloud services.

Key Takeaways

1. Canada ranks last in the G7 for publicly available compute, yet two-thirds of the federal government's $2 billion Sovereign AI Compute Strategy flows to private infrastructure rather than permanently governed public capacity.

2. Building data centres on Canadian soil doesn't guarantee digital sovereignty — without embedded governance, public compute investment will leak value to foreign technology giants.

3. Four models exist for sovereign compute: Crown corporation, Canadian-led P3, hyperscaler partnership, and fully private market — each offering different trade-offs between sovereignty, speed, cost, and governance.

4. Hyperscaler partnerships cannot guarantee Canadian legal jurisdiction over data; U.S. law, including the CLOUD Act, overrides any contractual sovereignty promises.

5. A layered architecture is the most viable path forward, matching different compute models to different use cases based on sovereignty requirements — with a publicly accessible, Canadian-anchored option as a non-negotiable component.

Publicly-Anchored Compute Infrastructure

A public option for cloud compute can take several forms, including:

- A Crown corporation with full government ownership,
- A Canadian-led public-private partnership (P3),
- Incumbent hyperscalers, operating with strong governance conditions, in a private cloud market under robust regulation.

Most of Canada's compute needs will be met by the market, including data centres located overseas or foreign-owned infrastructure located in Canada.

And that's appropriate; not every workload needs to be insulated from the influence of foreign hyperscalers.

However, a meaningful subset of Canada's compute capacity should be publicly anchored — accessible to the public and affordable enough to serve Canadian researchers, businesses, and institutions.

Publicly anchored compute will also operate exclusively governed by Canadian laws, norms, and values rather than those of a foreign jurisdiction.

"Publicly anchored" compute capacity does not exclusively mean publicly-owned infrastructure — other forms exist. Each option offers a different combination of sovereignty assurance, speed to capacity, cost to government, and governance depth.

Each option involves genuine trade-offs, and Canada, which currently ranks last in the G7 for publicly available compute infrastructure,[1] needs to make a deliberate and strategic choice rather than drifting into a default option.

This chapter examines four models for building a publicly anchored sovereign compute capacity in Canada, their benefits and downsides, and how they might serve or undermine Canadian sovereignty.

Why Canada Needs a Public Compute Option

The future of Canada's economy depends on access to compute. Compute is the ability to store, process and transfer data at scale, and is essential for business, public services, research, and national security. In all these cases compute is essentially the medium through which software and data are accessed and transformed to provide value to end users.

Compute is also essential for AI — both for training systems and for inference (using AI systems to process information). AI performance is closely tied to massive compute capacity, to the point that large scale compute has become a geopolitically important resource. The cost of hardware has doubled almost every year since 2019.[2]

In addition to AI's role in economic value creation, the technology is also becoming increasingly important in discovery and defence. Compute-heavy AI is being used to harden cybersecurity vulnerabilities.[3] AI will also likely be used for managing data-heavy systems like NORAD missile defence arrays and planned Arctic monitoring defence systems.[4]

The meteoric growth of the AI sector has coincided with geopolitical volatility that has led Canadian leaders and ordinary citizens to re-examine our sovereignty and national economic capacity.

In this context, Canada has already budgeted $2 billion towards the Sovereign AI Compute Strategy, in an effort to establish strategic compute capacity.[5]

Various policy advocates are calling on the government to do more.

However, as we have previously examined in the *Foundations of Digital Sovereignty* project, effectively building up sovereign capacity is not as simple as buying kit and setting up data centres on Canadian soil.

Peer countries have also recognized the need for sovereign, government-owned, public compute options. The United Kingdom, European Union member states, Japan, and India are all actively building publicly owned or publicly funded compute infrastructure as a strategic national asset.

As an example, the UK's AI Research Resource offers supercomputing access to researchers and SMEs for free.[6] France committed €109 billion to sovereign AI infrastructure,[7] and the EU has deployed a network of publicly backed "AI Factories" across the continent.[8] India's National AI Compute Initiative all reflects a similar imperative for government-mandated strategic compute.[9] Canada, which ranks last in the G7 for publicly available compute, is the outlier.[10]

Meanwhile, the Canadian federal government received criticism for courting 25-year cloud contracts exclusively with U.S. providers like Amazon, Google, Microsoft, Oracle after which they began considering Canadian companies as well.[11]

Nothing about the Canadian approach specifically addresses well-documented risks:

- Hyperscalers have consistently raised prices and pushed unwanted services onto clients;
- Over-reliance concentrates power in foreign hands;
- Exposure to the U.S. CLOUD Act means Canadian data is never fully beyond the reach of American law.

The Models for Sovereign Compute

The Government of Canada has recognized the compute gap and begun to act, but the design of our response thus far suggests limited ambition. The Canadian Sovereign AI Compute Strategy, backed by a combined $2 billion across Budgets 2024 and 2025, is the most significant federal compute investment in Canadian history.[12] The headline figure, however, obscures a more modest reality.

The strategy has three components:

- Up to $700 million through the AI Compute Challenge to mobilize private commercial data centre investment;
- Up to $1 billion for public supercomputing infrastructure including through the Sovereign Compute Infrastructure Program (SCIP);
- $300 million in the AI Compute Access Fund to subsidize innovator access.[13]

Two-thirds of the Sovereign AI Compute Strategy funding flows toward private infrastructure or access subsidies, not toward publicly owned, permanently governed sovereign capacity.

Notably, Canada's existing sovereign compute strategy is implicitly geared to a series of different smaller projects, to spread the money between researchers and businesses. This approach is likely to fail, as data centres benefit from scale. One big megacompute project would likely deliver greater benefits.

The SCIP in particular has notable weaknesses. Only consortia led by a non-profit or post-secondary institution incorporated in Canada may apply — private firms may participate as partners but cannot lead.[14]

This effectively means rooting significant compute infrastructure within universities, rather than leveraging private sector experience in infrastructure development.

SCIP is a grant program to build a supercomputer but doesn't build the institution needed to own and govern that supercomputer for decades. Canada has previously built publicly funded research infrastructure, and we have then watched it atrophy when the funding cycle ended.

Without permanent governance architecture built into the institution itself — a mandate, accountability mechanisms, or a long-run operational model — the Sovereign AI Compute Strategy risks the same outcome: capable infrastructure built at public expense, sovereign in name, but without the institutional permanence to remain so. What's worse is the many different projects being funded across the Canadian Sovereign AI Compute Strategy creates a limited ability to infuse Canadian standards and Canadian governance in our sovereign compute capacity.

We can imagine several structures for sovereign Canadian compute which would allow for better governance and an increased ability to strategically deploy compute power for Canadian national priorities.

For *Foundations of Digital Sovereignty* we have looked at three main options:

- A Crown corporation with full government ownership;
- A Canadian-led public-private partnership (P3);
- Incumbent hyperscalers, operating with strong governance conditions, in a private cloud market operating under robust regulation.

Option 1: Crown Corporation Model

A Crown corporation mandated to build and operate sovereign Canadian compute capacity at scale is arguably the strongest guarantee of sovereignty — full legal jurisdiction, genuine operational control, governance being embedded from day one.

However, a Crown corporation would be slowest to build, the most institutionally demanding, and there is a track record in Canada of other Crown corporations being financially unstable over time.[15]

A Crown corporation would have the opportunity to embed governance structures and have strategically adopted Canadian standards from the beginning.

The entity could begin with steady funding from the government, but then transition into a self-sustaining institution — which could potentially even operate as a for-profit entity by charging for use of sovereign Canadian compute.[16]

A Crown corporation would also be positioned to provide compute capacity to researchers and SMEs at stable and affordable rates. The Crown corporation could be insulated from market competition and price fluctuations if primarily funded by the Canadian government.

It could also operate as a competitive and for-profit entity,[17] driving the market prices down and increasing accessibility.

However, there are also significant downsides to Crown corporations that cannot be ignored. Crown corporations are slower to capitalize and build than private alternatives, and Canada cannot wait a decade for procurement cycles and enabling legislation to produce operational capacity.

The political economy is difficult: Crown corporations, especially at a nascent stage, are vulnerable to budget cycles as well as political interference.

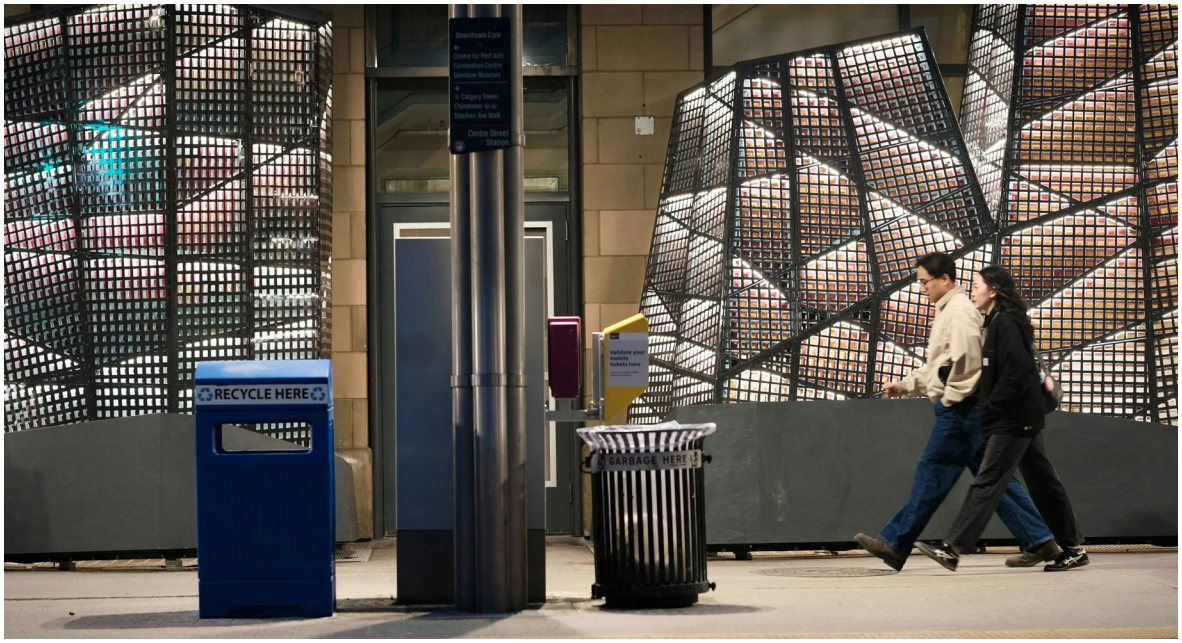

In technology-adjacent sectors, Canadian Crown corporations also have a mixed record: Atomic Energy of Canada Limited is a genuine success story, but the sector has also produced cautionary tales of procurement failures, cost overruns, and institutions that outlived their mandate.

The most underappreciated risk, however, is crowding out the emerging private sector Canadian compute providers. Emerging compute infrastructure builders that are genuinely Canadian, such as Telus, Micrologic and eStrutxure, need a competitive market to grow into. A Crown corporation with preferred procurement status distorts that market, which can potentially stunt the domestic private capacity that Canada also needs to build.

Finally, attracting and retaining the talent needed to run frontier chip clusters is genuinely difficult within public sector compensation structures.

Option 2: Hyperscaler-Led Public-Private Partnership

Without a significant change in direction from the federal government, the default option for larger projects will be some sort of long-term public-private partnership with foreign hyperscaler cloud service providers who can provide the scale of compute the government seeks to make available.

This is the model implicit in the federal government's 25-year cloud contracting discussions with U.S. providers, and it is the path of least resistance given how far behind Canada currently sits in domestic compute capacity.

The promise of the hyperscaler-led P3 model for public compute is that these companies can deliver capacity quickly.

Hyperscalers have the ability to quickly deliver what Canada urgently needs: frontier chip clusters, like the NVIDIA Vera Rubin and Grace Blackwell architecture that train and run competitive AI models, deployed at scale, with global redundancy and mature operational infrastructure.[18]

Microsoft's announced $19 billion plan to expand AI and cloud infrastructure in Canada already includes new data centre projects coming online in 2026 and an increase in compute capacity nationwide.[19]

For proponents, this represents real capacity delivered immediately, at no direct capital cost to the government, with contractual sovereignty protections attached.

As part of this push, Microsoft is committing to keep Canadian data on Canadian soil, and writing into contracts that Microsoft will challenge foreign government demands for Canadian data where it has the legal grounds to do so.

However, as we have seen in earlier chapters of *Foundations of Digital Sovereignty*, the value of these commitments is dubious. Although a government agreement could enshrine contractual obligations for preserving sovereignty, these would not be able to assure sovereignty and security in the face of U.S. laws.

In June 2025 testimony before a French Senate committee, Microsoft France's Director of Public and Legal Affairs was asked under oath whether he could guarantee that data could not be transmitted to the U.S. government without French approval. His answer was: "No, I cannot guarantee that, but again, it has never happened before."[20]

Microsoft's promise to "rigorously defend the uninterrupted operation of cloud services for Canadian government customers" and fight any order to suspend operations in Canada implicitly concedes that such orders are legally possible.[21] A promise to resist using kill switches is not the same as escaping the kill switch entirely.

The bottom line is that convenience and expediency elides the risks to Canada's digital sovereignty. If the government is essentially only acting as a customer to enormous foreign corporations, we are running up against a significant power imbalance in any governance that Canada attempts to impose.

While we may try to dictate Canadian standards, or Canadian data usage rules, the vendors would be able to respond with a take-it-or-leave-it approach. And if Canada's starting point in this exercise is prioritizing convenience and expedience over real governance, it's hard to imagine that we would be willing to walk away.

The hyperscaler partnership model has a legitimate role for commercial compute workloads, capacity, and use cases where data sensitivity is low and speed of access is the dominant requirement. But it's structurally insufficient for any workload where Canadian legal jurisdiction must be guaranteed: defence-adjacent systems, sensitive government data, and research involving personal or national security information.

In an environment where U.S. belligerence is becoming increasingly worrying for governments and businesses, organizations might decide that even non-private, non-security essential tasks might need to be done on non-hyperscaler infrastructure.

A major risk with the hyperscaler P3 approach is that it further entrenches their market dominance, in effect stifling nascent Canadian alternatives.

Using it as the primary model for sovereign compute would mean our sovereignty rests on the contractual goodwill of American corporations operating under American law — which is precisely the dependency this chapter argues Canada must reduce.

Option 3: Canadian Public-Private Partnership

A Canadian P3 is the most promising middle path — faster and cheaper than a Crown corp., more governable than a hyperscaler deal — but only if the governance conditions attached to public funding are binding and enforceable, which the current Sovereign Compute Infrastructure Program does not guarantee.

Unlike a Crown corporation, under a P3 model the government would not own the infrastructure outright; instead, it would use funding agreements, contractual obligations, and regulatory conditions to steer privately or institutionally operated systems toward public interest outcomes. This is broadly the model SCIP is already focused on for building compute specifically for research purposes. However, expanding Canadian P3 options to also try to address broader compute needs, rather than partnering with U.S. hyperscalers, could lead to greater compute sovereignty.

Proponents argue that co-investment leverages private capital and operational expertise, which the government does not have in-house, accelerating build timelines and reducing direct fiscal exposure.[22]

Canada has demonstrated capacity for this model in adjacent sectors: ENCQOR 5G brought together five global technology leaders including one Canadian firm — Ericsson, Ciena, IBM, Thales, and CGI — with federal and provincial governments in a $400 million partnership, and public funds that were matched dollar-for-dollar by the private sector.[23]

The result was real shared infrastructure: a network of 5G testbeds giving Canadian SMEs and researchers access to a state-of-the-art development and testing platform. A research-based partnership like this also distributes risk if the technology or market conditions shift.

However, the structural limitations are significant, and they matter more for compute sovereignty than they did for 5G testbeds. ENCQOR worked because it was an innovation access model: a P3 designed to enable SMEs and researchers to test and scale pre-commercialized applications.[24] It was not a permanent sovereign infrastructure play and it relied on several foreign multinationals alongside the Canadian firm CGI. The anchor firms retained control of their underlying technology throughout.

In the case of sovereign compute, the partnership would need to produce infrastructure that is Canadian-governed permanently, operationally independent from foreign legal regimes, and accountable to public interest mandates over decades. This would be far beyond the scope of what ENCQOR was designed or tested to deliver.

On longer timeline, Canadian partner firms that scale globally would potentially want to leave the consortium, build their own infrastructure if they aim to scale, or they could be acquired by foreign firms, creating a backdoor for non-sovereign actors.

There are also long-running critiques of the P3 model that would be applicable to this arrangement.

P3s tend to cost taxpayers more over the long run, transfer less risk than advertised, and produce accountability gaps that are difficult to close after contracts are signed.

A Canadian P3 for compute also faces a partner problem — the domestic firms capable of building at the required scale are few, meaning the natural co-investors are either U.S. hyperscalers or foreign hardware vendors, both of which reintroduce sovereignty exposure the Canadian-focused partnership was meant to solve.

Furthermore, the few Canadian firms that are eligible for such a project are likely to be subject to U.S. jurisdiction themselves through the CLOUD Act — any Canadian company with enough of a presence in the U.S. is captured.[25]

These risks can be mitigated to some extent with appropriate governance, for example through robust contractual requirements or the modernization of Canada's blocking statute, as proposed earlier.

All challenges aside, such a partnership might be very successful if the government chooses to commit to the approach and seeks to secure private capital and funding from Canadian institutional funds.

Option 4: Fully Private, Market-Led, and Regulated

A fully private model might be appropriate for what Canadian businesses and consumers need from compute most of the time.

However, the Cohere case illustrates its structural limit: when Canadian firms lack the scale to build independently, 'private Canadian' rapidly transforms into 'private American' with some sovereignty washing to make it seem acceptable.

Proponents of a completely private market approach might argue that government ownership and partnership arrangements crowd out private investment, distort market signals, and produce the kind of procurement inefficiencies and cost overruns that have plagued public infrastructure projects across multiple sectors.

The private sector moves faster, operates leaner, and has stronger incentives to keep pace with rapidly evolving hardware generations — a genuine advantage in a sector where the frontier shifts every 18 months.

Canada already has a version of this logic at work: the federal government's $240 million investment in Cohere's $725 million project to expand domestic compute capacity in Canada was explicitly framed as crowding in private capital, with the government acting as anchor investor rather than owner.[26]

Taking a hands-off approach and letting the private market handle everything is also a more attractive option for the people who argue that we don't really need to worry about foreign interference under the CLOUD Act and FISA.

We would argue that this position is shortsighted and perpetuates the same kinds of mistakes about technology governance that Canada has been making for decades, in allowing the complete capture and foreign control of our digital infrastructure.

But the critical weakness of this model surfaces in the Cohere case itself.

Cohere used the federal funding to partner with CoreWeave, a U.S. AI infrastructure company, to build the $725 million data centre project.

The government contract contained no requirement to choose a Canadian supplier, so the money ultimately went to an American firm to build the infrastructure.[27] This is the structural problem with market-led compute. We are dealing with a market that American companies dominate. Without a clear and coherent strategy, foreign firms will maintain that dominance and Canadian companies will offer 'sovereignty' as a service but actually ultimately serve as a backdoor for foreign interests.

Building data centres locally doesn't guarantee digital sovereignty if those facilities are run by foreign entities.

Private Canadian firms competing at the frontier will consistently face pressure to partner with U.S. firms that have preferential access to Nvidia chips, CoreWeave infrastructure, and U.S. capital markets.

Regulation can set residency and governance floors, but it cannot manufacture Canadian-owned capacity that does not yet exist, and it can't insulate privately owned infrastructure from foreign legal regimes any more than contractual promises can.

A Sovereign Compute Strategy That Works

The four models examined here are not equally suited to every compute need, and Canada is unlikely to end up with a pure version of any one of them. The most sensible outcome is a layered architecture: different models applied to different use cases based on the sovereignty requirements, security classification, and public interest obligations each context demands.

The fully private model might be appropriate for commercial compute. The hyperscaler partnership could have a legitimate role for workloads where sovereignty requirements are limited and speed is the dominant concern.

The Canadian P3 is the most promising middle path if, and only if, the governance conditions attached to public funding are binding, enforceable, and designed to produce permanent Canadian-anchored infrastructure instead of time-limited access.

The Crown corporation remains the strongest sovereignty instrument for the workloads where legal jurisdiction, operational control, and embedded governance are non-negotiable — and those workloads exist, and will grow.

In assessing what option, or what mix of options Canada should weave together for our compute needs, there will be a need to take a variable geometry approach as mentioned by Prime Minister Carney and expanded upon in great detail for the AI context by Professor Mark Daley.[28]

Assessing how dual-use, how secret and secure, and how strategically important a task or technology is will be essential in determining what the best compute supplier and associated arrangement might be.

But ultimately some form of publicly accessible option will be necessary.

What the analysis makes clear is that structural choice for Canadian sovereign compute can't be separated from the prior question of what sovereignty requires.

The federal government's current approach defines sovereignty primarily in terms of locating infrastructure in Canada, prioritizes speed over institutional permanence, and routes most public investment through private and P3 channels.

That may be the right set of trade-offs for some compute needs. However, it is not sufficient for all of them. The program as designed does not create the institutional architecture to govern the infrastructure it is building over the long run.

Solutions

In order to build the necessary compute to meet Canada's strategic needs, there will need to be a structured and principled effort to determine what those needs are and what model suits it best. The government of Canada can do this by:

1. Laying out a clear strategic plan on what areas of compute will be mission critical, choosing domains and areas of research that we are seeking to lead on.
2. Mapping that strategy around security, privacy, and sovereignty needs to determine what model of public compute is best.
3. Choosing between either a Crown corporation or a Canadian-led P3 option, or potentially a mix of both, that can meet the strategic needs outlined in the strategy and sovereignty mapping.

Compute infrastructure is not the foundation of Canada's digital sovereignty — governance is. But governance without infrastructure is incomplete. A country that sets excellent rules over systems it doesn't own, cannot access affordably, and cannot insulate from foreign legal regimes will find those rules tested at precisely the moments they matter most.

References

1 Graham Dobbs and Jake Hirsch-Allen, "Can Canada Compute? Policy Options to Close Canada's AI Compute Gap," The Dais, March 2024, https://dais.ca/reports/can-canada-compute/. Updated data as of November 2025 obtained from the original source cited in "Can Canada Compute," taken from, Top500, "List Statistics," accessed June 1, 2026, https://www.top500.org/statistics/list/.

2 Konstantin F. Pilz, et al., "Acquisition costs of leading AI supercomputers have doubled every 13 months," Epoch AI, June 5, 2025, https://epoch.ai/data-insights/ai-supercomputers-cost-trend.

3 Chris Rohlf, "AI and the Software Vulnerability Lifecycle," Center for Security and Emerging Technology, August 4, 2025, https://cset.georgetown.edu/article/ai-and-the-software-vulnerability-lifecycle/.

4 Department of National Defence, "Annex C: Canada's NORAD Modernization Plan," Government of Canada, last modified April 17, 2024, https://www.canada.ca/en/department-national-defence/corporate/reports-publications/north-strong-free-2024/annex-c-canada-norad-modernization-plan.html.

5 Innovation, Science and Economic Development Canada, "Canadian Sovereign AI Compute Strategy," Government of Canada, last modified October 31, 2025, https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy.

6 Department for Science, Innovation and Technology and UK Research and Innovation, "AI Research Resource," Gov.UK, last modified November 7, 2025, https://www.gov.uk/government/publications/ai-research-resource.

7 Élysée, "Make France an AI powerhouse," February 11, 2025, https://www.elysee.fr/en/emmanuel-macron/2025/02/11/make-france-an-ai-powerhouse.

8 European Commission, "AI Factories," last modified April 23, 2026, https://digital-strategy.ec.europa.eu/en/policies/ai-factories.

9 IndiaAI, "IndiaAI Compute Capacity," accessed May 29, 2026, https://indiaai.gov.in/hub/indiaai-compute-capacity.

10 Dobbs and Hirsch-Allen, "Can Canada Compute?" Updated data as of November 2025 from Top500, "List Statistics."

11 Bill Curry, "Ottawa to work with Canadian cloud providers after industry pushback over U.S. shortlist," The Globe and Mail, April 21, 2025, https://www.theglobeandmail.com/politics/article-ottawa-to-favour-canadian-cloud-providers-after-industry-pushback-over/.

12 Innovation, Science and Economic Development Canada, "Canadian Sovereign AI Compute Strategy;" Government of Canada, *Canada Strong: Budget 2025*, November 2025, https://budget.canada.ca/2025/report-rapport/pdf/budget-2025.pdf.

13 Innovation, Science and Economic Development Canada, "Canadian Sovereign AI Compute Strategy."

14 Innovation, Science and Economic Development Canada, "AI Sovereign Compute Infrastructure Program," Government of Canada, last modified June 1, 2026, https://ised-isde.canada.ca/site/ised/en/ai-sovereign-compute-infrastructure-program.

15 Glen Hodgson, "What To Do with Canada Post?," C.D. Howe Institute, June 6, 2024, https://cdhowe.org/publication/glen-hodgson-what-do-canada-post/.

16 Treasury Board of Canada Secretariat, "Directors of Crown corporations: an introductory guide to their roles and responsibilities," Government of Canada, last modified October 28, 2025, https://www.canada.ca/en/treasury-board-secretariat/services/guidance-crown-corporations/directors-crown-corporations-introductory-guide-roles-responsibilities.html.

17 Treasury Board of Canada Secretariat, "Directors of Crown corporations."

18 NVIDIA, "Data Center Products," accessed June 3, 2026, https://www.nvidia.com/en-us/data-center/products/.

19 Brad Smith, "Microsoft Deepens Its Commitment to Canada with Landmark $19B AI Investment," Microsoft, December 9, 2025, https://blogs.microsoft.com/on-the-issues/2025/12/09/microsoft-deepens-its-commitment-to-canada-with-landmark-19b-ai-investment/.

20 Senate of France, "Hearing of Mr. Anton Carniaux, Director of Public and Legal Affairs, and Mr. Pierre Lagarde, Technical Director for the Public Sector, of Microsoft France," June 10, 2025, https://www.senat.fr/compte-rendu-commissions/20250609/ce_commande_publique.html#toc2

21 Smith, "Microsoft Deepens Its Commitment to Canada."

22 The Canadian Council for Public-Private Partnerships, "What are public-private partnerships (P3s)?," accessed May 29, 2026, https://www.pppcouncil.ca/why-p3s/what-are-p3s; The Canadian Council for Public-Private Partnerships, "Types of P3s," accessed May 29, 2026, https://www.pppcouncil.ca/why-p3s/types-of-p3s.

23 ENCQOR, "ENCQOR 5G Mission Accomplished: Final Activity Report", March 31, 2023, https://encqor.ca/wp-content/uploads/2024/06/23_12-Rap-Annuel_encqor5g_ENG-19DEC.pdf.

24 ENCQOR, "Final Activity Report".

25 Barry Appleton, "Whose Law Governs Canadian Data? The CLOUD Act, Executive Agreements, and Digital Sovereignty," SSRN (Working Paper), December 22, 20256, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5955017.

26 Department of Finance, "Deputy Prime Minister announces $240 million for Cohere to scale-up AI compute capacity," Government of Canada, December 4, 2024, https://www.canada.ca/en/department-finance/news/2024/12/deputy-prime-minister-announces-240-million-for-cohere-to-scale-up-ai-compute-capacity.html.

27 The Canadian Shield Institute for Public Policy, "Sovereignty Score: Cohere Investment", November 19, 2025, https://img1.wsimg.com/blobby/go/e37fd200-232f-4959-9dca-2108327c2abf/downloads/df90526f-3ef5-4182-8c82-1d0e889dc3fd/The%20Cohere%20Investment.pdf?ver=1775754443107.

28 Mark Daley, "Carney's Variable Geometry Needs Constraints in the AI Age," Centre for International Governance Innovation, May 7, 2026, https://www.cigionline.org/articles/carneys-variable-geometry-needs-constraints-in-the-ai-age/.