Sovereign Compute Infrastructure

Underwriting National AI Capacity as an Asset Class

Authored by Chennakeshav Adya
Independent Researcher
August 2026

Abstract

Governments are treating advanced artificial-intelligence compute as research infrastructure, industrial capacity, a security capability and a platform for public services. Canada is funding sovereign infrastructure and compute access. The European Union is building AI Factories around EuroHPC systems. The United Kingdom is expanding the AI Research Resource. India operates an empanelled-provider compute portal with published access and pricing mechanics. The United States National Artificial Intelligence Research Resource combines public and contributed resources. Singapore links new data-centre capacity to economic contribution, resilience and sustainability.

These programmes show that national compute cannot be underwritten as a conventional data-centre shell with servers added. Compute value depends on a complete service chain: dependable power and cooling, accelerators and networks, secure data and storage, software environments, workload scheduling, user support, eligibility and allocation, model governance, cyber security and technology refresh. A technically impressive system can produce weak economic value when users cannot access it, workloads are poorly matched, utilisation is unmeasured or refresh funding is absent.

This paper develops a decision framework for underwriting sovereign compute infrastructure as an asset class. It defines the control perimeter, segments demand, maps the service stack, distinguishes delivery models, tests utilisation and refresh economics, allocates risks across public and private participants, and sets a 180-day development sequence. Six figures present the sovereign control perimeter, demand ladder, delivery-model spectrum, cash-flow architecture, refresh cycle and investment gate. Six tables provide an international programme comparison, diligence file, delivery-model matrix, hypothetical financing case, stress matrix and execution plan.

Every price, capacity, utilisation rate, cost, subsidy, payment, schedule and financial result in the worked case is a hypothetical management assumption created solely to demonstrate the method. The paper does not claim an actual project, contract, procurement, appropriation or investment return.

Sovereignty, national security, data protection, public procurement, state aid, foreign investment, export control, sanctions, competition, cyber security, energy, planning, environmental, financing, accounting, tax and investment decisions require advice from qualified professionals in the relevant jurisdictions. This paper provides general information for professional audiences and does not provide legal, regulatory, security, engineering, tax, accounting, valuation, credit or investment advice.

JEL Classification: G31, G32, H54, L86, O32, O38, Q40

Keywords: sovereign compute, artificial intelligence, data centres, digital infrastructure, project finance, public investment, compute access, accelerator refresh, national infrastructure

1. Define the sovereign outcome

A national compute programme should begin with the outcome that requires public or sovereign control. Possible outcomes include advanced research, strategic model development, protected public-sector workloads, affordable access for domestic firms, scientific discovery, workforce development, industrial adoption and resilience when external supply is constrained.

These outcomes require different infrastructure. Frontier model training needs large, tightly connected accelerator clusters and specialised engineering. Scientific users may need diverse architectures, high-performance storage and long runs. Small firms may value reserved cloud capacity, simple onboarding and technical support. Protected government workloads may require separate security domains, accredited operations and strong control over management interfaces.

Sovereignty is therefore a set of enforceable decision rights. It can cover the physical location of equipment, ownership, board control, operating authority, data residency, cryptographic keys, identity systems, software, management planes, supply-chain decisions, workload allocation, pricing, refresh and termination.

Canada's 2026 Sovereign Compute Infrastructure Program guide requires sustained Canadian control across design, build, operation and refresh, and explicitly considers data, metadata, access pathways, control planes and monitoring tools.[1] This shows why a local building alone is insufficient.

The investment committee should approve the required sovereign outcomes and their control perimeter before selecting a technology or financing structure. Every additional control objective can affect cost, supplier choice, scale, interoperability and utilisation.

Figure 1. Sovereign control perimeter

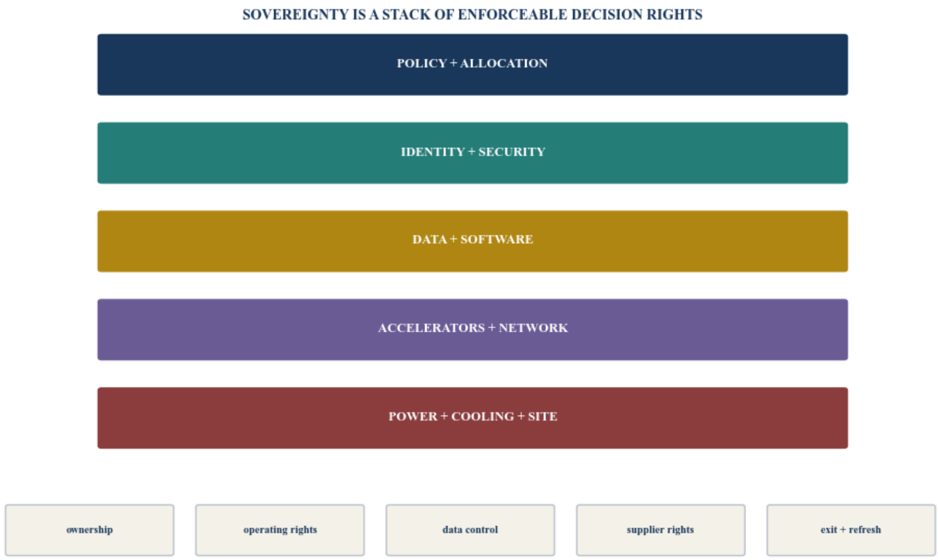

Author framework. The required control level should be approved for every layer rather than assumed from physical location.

2. Treat compute as a service chain

The productive asset extends beyond accelerators. A complete service needs power, heat rejection, network fabric, storage, data movement, orchestration, schedulers, development environments, models, cyber controls, user support, metering and billing.

The weakest layer can strand the whole investment. A cluster without adequate network bandwidth can leave accelerators waiting. Storage bottlenecks can slow training. Incomplete software images can delay onboarding. Complex allocation rules can suppress demand. Missing support can prevent smaller organisations from converting accelerator hours into usable outcomes.

The European Union describes AI Factories as hubs built around AI-optimised supercomputers with computing, data and support services for research, start-ups, small and medium enterprises and industry.[2][3] The United States NAIRR pilot similarly combines computation, data, software, models, training and user support.[4] These programmes frame national capacity as a service ecosystem.

Underwriting should therefore allocate capital and operating budget across the full chain. Accelerator procurement may dominate initial equipment cost, while software, security, data, people and refresh determine continued service.

3. Segment demand before sizing supply

National demand is heterogeneous. It should be divided by user, workload, security, duration, scale and willingness to pay. Research, start-up development, regulated-industry deployment, public administration and strategic model training should not share one undifferentiated forecast.

Each segment requires a workload envelope. Important variables include accelerator type, memory, interconnect, storage throughput, data sensitivity, run duration, queue tolerance, software environment, support intensity and portability.

Evidence of demand can progress from expression of interest to benchmarked workload, approved allocation, budgeted service order, take-or-pay capacity and contracted minimum use. These stages should carry different probability weights.

The United Kingdom AIRR offers access routes through Isambard-AI and Dawn for researchers, academia and industry.[5] IndiaAI uses eligibility, project review, bills of materials, priced services and approval thresholds on its compute portal.[6][7] These mechanisms reveal demand through actual applications and resource requests rather than broad policy interest.

Figure 2. Demand evidence ladder

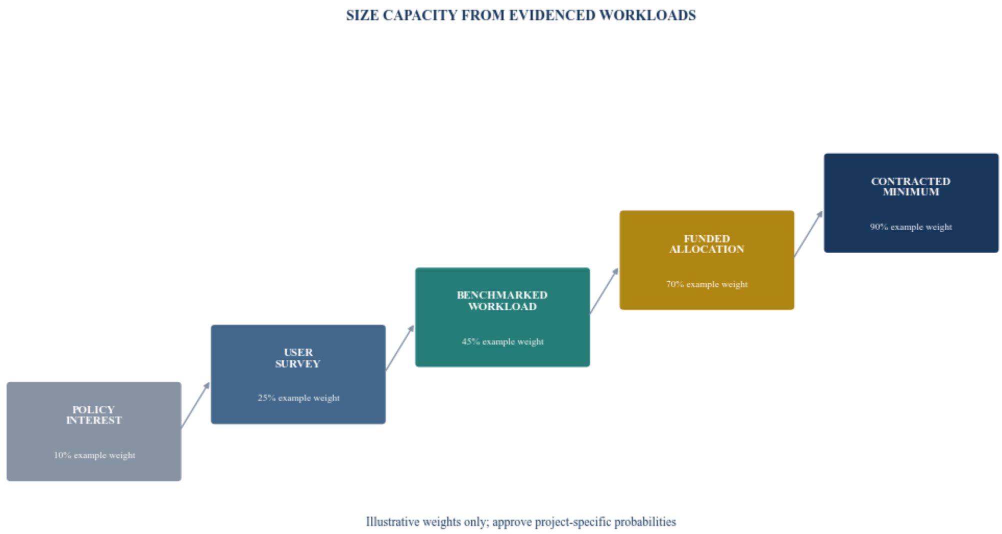

Author framework. Forecast weight should rise as users provide technical, budget and contractual evidence.

4. Read international programmes as design evidence

Countries are combining direct ownership, grants, access subsidies, commercial panels, federated systems and contributed resources. Their programmes should be read within their stated purposes and jurisdictions.

Canada announced up to CAD 2 billion for its sovereign compute strategy in December 2024.[8] Its later infrastructure programme described an approximately CAD 890 million infrastructure-build layer, subject to appropriations and formal agreements.[9] The separate Compute Access Fund supports eligible firms' use of Canadian and, for a limited period, non-Canadian services under stated contribution rules.[10]

The European Union is building AI Factories around EuroHPC supercomputers and connecting computing with data, support and industrial services.[2][3] The UK is expanding AIRR through public supercomputers and cloud procurement, with differentiated access routes.[5][11]

IndiaAI aggregates commercial providers, publishes configurations and prices, applies subsidy rules and routes users through a national portal.[6][7] The United States NAIRR pilot federates public and private contributions across agencies and resource providers.[4][12] Singapore's second data-centre call evaluates capacity proposals for AI contribution, connectivity, resilience, economic value and sustainability.[13]

Table 1. International programme design comparison

Programme	Infrastructure model	Access mechanism	Underwriting lesson
Canada sovereign compute	contribution-funded domestic infrastructure and public capacity	programme agreements and access support	separate build capital, service layer, appropriation and user economics
European AI Factories	AI-optimised EuroHPC systems with data and support services	factory calls and user programmes	value depends on ecosystem services and cross-site coordination
United Kingdom AIRR	public supercomputers plus planned cloud integration	gateway, rapid, open and strategic routes	match allocation route and service level to workload class
IndiaAI Compute	panel of empanelled cloud providers	portal, eligibility, bill of materials, pricing and subsidy	transparent configurations and metering can aggregate distributed supply
United States NAIRR	public resources plus private contributions	national pilot allocations and open resources	federation can widen capacity while increasing contract and continuity complexity
Singapore data-centre call	competitive allocation of scarce capacity	assessed proposals from private developers	grid, sustainability, resilience and economic additionality are part of capacity value

Official programme evidence is summarised within each programme's stated mandate and jurisdiction.

5. Select the delivery model

Five broad structures are available. Direct public ownership can maximise control and public access while placing technology, utilisation and refresh risk on the state. A concession or availability-payment model can transfer construction and service obligations while retaining public allocation rights. A joint venture can align government, operator, utility and technology participants but requires clear governance.

Reserved commercial capacity can accelerate access and preserve flexibility. It can also create dependence on provider terms, control planes and foreign jurisdictions. A federated marketplace can aggregate multiple providers and architectures, though service consistency and continuity become more complex.

The appropriate structure depends on sovereign control, speed, demand certainty, workload scale, security, procurement authority, public accounting, financing capacity and desired competition. A programme may use several structures simultaneously. Large protected workloads could use a dedicated system while smaller firms access a commercial panel.

Figure 3. Delivery-model spectrum

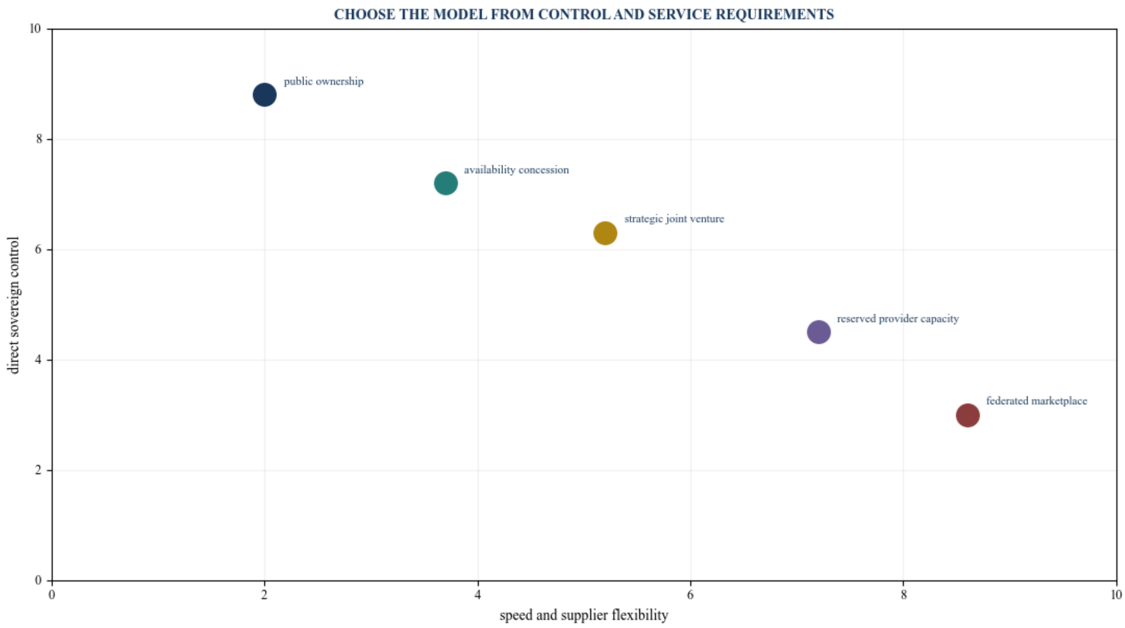

Author framework. Control, speed and risk transfer vary by project and contract.

6. Establish the asset boundary

The project company may own land, building, electrical and cooling plant, network equipment, accelerators and software. Other structures separate durable facility assets from short-lived technology and the national service layer.

This separation matters because economic lives differ. A power connection, substation and data hall may support several technology generations. Accelerators can lose economic relevance much faster. Proprietary platform software can create switching cost without durable resale value.

The asset register should state owner, location, title, useful life, warranty, support, licence, security classification, export restriction, replacement duty, permitted use, insurance and residual-value assumption for every material component.

Financing should follow controllable cash and collateral. Long-life infrastructure may support longer-tenor debt when power, site and service payments are dependable. Accelerator debt needs shorter amortisation, stronger purchase and support rights, conservative residual value and a funded refresh mechanism.

7. Secure power, cooling and network capacity

Accelerator capacity is unusable without dependable electricity, heat rejection and network connectivity. The underwriting file should distinguish utility connection, firm import capacity, on-site generation, backup systems, renewable supply, curtailment rights and fuel continuity.

Power-to-compute conversion varies with hardware utilisation and facility overhead. A nominal megawatt figure should be translated into supported racks, usable accelerators and customer-serving hours under normal, hottest-day and degraded conditions.

Cooling design should match rack density, climate, water availability, equipment warranties and redundancy. Network design should cover internal fabric, storage, external connectivity, route diversity, latency, cyber control and data-transfer cost.

Singapore's 2025 data-centre call linked new capacity to resilience, international connectivity, economic value and sustainability, with at least 200 MW stated for the call.[13] This demonstrates that scarce power allocation can be conditional on national value rather than treated as a passive utility input.

8. Underwrite the accelerator portfolio

A sovereign platform may contain several accelerator generations and vendors. Portfolio design can reduce supply dependence and match workloads, while heterogeneous environments add software, scheduling and support complexity.

Underwriting should use supported workload performance rather than peak theoretical computation. The evidence file should include benchmark results, memory capacity, interconnect, power, cooling, software compatibility, availability, support, lead time, export permissions and delivered price.

New technology introduces price-performance upside and deployment risk. Mature systems provide better operating evidence but can face faster displacement. Procurement may therefore use staged commitments, acceptance tests, price adjustment, optional expansions and refresh windows.

Residual value should remain conservative. Secondary-market price, export restrictions, support status, removal cost, data destruction and energy efficiency affect recoverability. A zero or low base residual can be appropriate when debt repayment depends primarily on service cash.

9. Build the service layer

Users require an accessible environment. The service layer should include identity, project approval, quota, scheduler, software images, storage, data movement, monitoring, security, support, incident management, metering and billing.

IndiaAI's portal publishes eligible user categories, project and approval mechanics, provider configurations, on-demand and reserved pricing, storage and potential subsidy calculations.[6][7] The UK AIRR publishes access routes for different users and scales.[5] These features turn infrastructure into an allocable service.

Service-layer performance should be measured through onboarding time, job start time, completed workload, failed-job rate, accelerator utilisation, support response, security events, user concentration and cost per useful accelerator hour.

The operator should maintain portable interfaces and export paths. Users should be able to retrieve data, models, logs and artefacts in documented formats. Platform convenience should not become an uncontrolled national dependency.

10. Design allocation and pricing

Allocation determines who receives scarce capacity, on what terms and for which outcomes. It should be transparent enough for accountability and flexible enough to respond to strategic priorities.

Possible mechanisms include peer-reviewed research allocation, ministerial or agency priority, competitive industrial calls, priced commercial access, vouchers, matched funding, reserved public capacity and emergency override. Each route needs eligibility, decision rights, conflicts management, usage expiry and reporting.

Pricing can recover full cost, marginal cost, a policy-set tariff or a subsidised share. A programme can publish rates while applying grants or credits separately. This preserves visibility into the economic value of capacity.

Subsidy should have a defined objective and counterfactual. It can target early-stage firms, national-priority science, protected workloads, capability building or domestic commercialisation. The programme should measure whether supported users complete workloads and achieve stated outputs rather than counting allocated hours alone.

Table 2. Compute access and pricing file

Decision	Evidence required	Contract or rule	Operating measure
Eligibility	legal status, user class, project purpose and security	approved categories and exclusions	applications and acceptance rate
Technical allocation	workload benchmark, duration, architecture and data	quota, queue priority and expiry	requested versus completed hours
Price	service cost, market comparator and policy objective	tariff, discount, credit and reset	realised revenue and subsidy per hour
Data control	classification, residency, access and retention	approved environment and export rules	incidents, exceptions and deletion evidence
Support	user maturity and workload complexity	included service and escalation	onboarding time and support effort
Outcome	research, product, service or public benefit	reporting and milestone duties	completed artefacts and adoption

The file should connect every allocation to authority, technical need, price, control and measured outcome.

11. Separate revenue from public benefit

The cash model may include user charges, capacity subscriptions, government service payments, grants, appropriations, research funding, anchor-tenant payments and platform fees. These sources should remain separate.

User charges can demonstrate demand but may not cover the public-service capacity reserved for research, security or emergency needs. Availability payments can support financing when they are appropriated, contractually clear and linked to measurable service. Grants reduce capital but do not fund perpetual refresh unless expressly designed to do so.

Strategic benefits such as talent, research leadership, national resilience and retained intellectual property can support policy approval. They should not be inserted into project cash flow unless they generate an enforceable payment.

The base financing case should use contracted or otherwise dependable cash. Subsidised demand should be traced to the funding source and duration. Upside from commercialisation, ecosystem growth or future model demand should remain a scenario until supported.

Figure 4. Sovereign compute cash-flow architecture

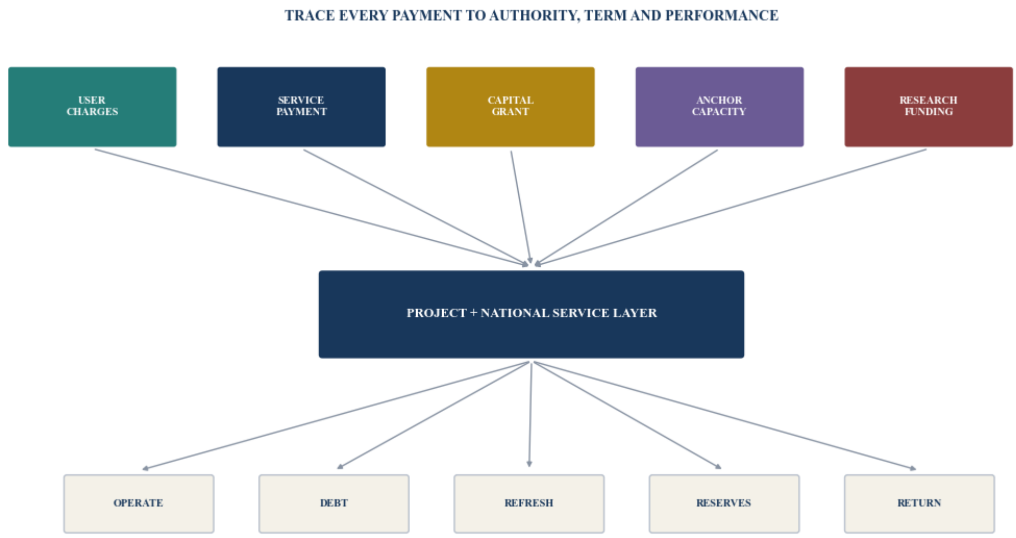

Author framework. Public benefit and project cash should be measured separately.

12. Fund the refresh cycle

Compute infrastructure requires recurrent technology renewal. The initial capital decision should include the timing, scope, procurement authority and funding source for refresh.

A refresh reserve can be funded from service payments, user revenue or appropriations. It should be held under clear rules and sized from current replacement evidence, installation cost, software migration, decommissioning and downtime.

Staggered replacement can reduce a single cliff and maintain service diversity. It can also create a heterogeneous platform that is harder to support. Whole-cluster replacement can preserve consistent performance while concentrating capital and execution risk.

Contract design should address successor technology, benchmark equivalence, vendor substitution, price movement, export restrictions and software migration. A fixed specification for a distant delivery date can lock the programme into outdated equipment.

Figure 5. Technology refresh and financing cycle

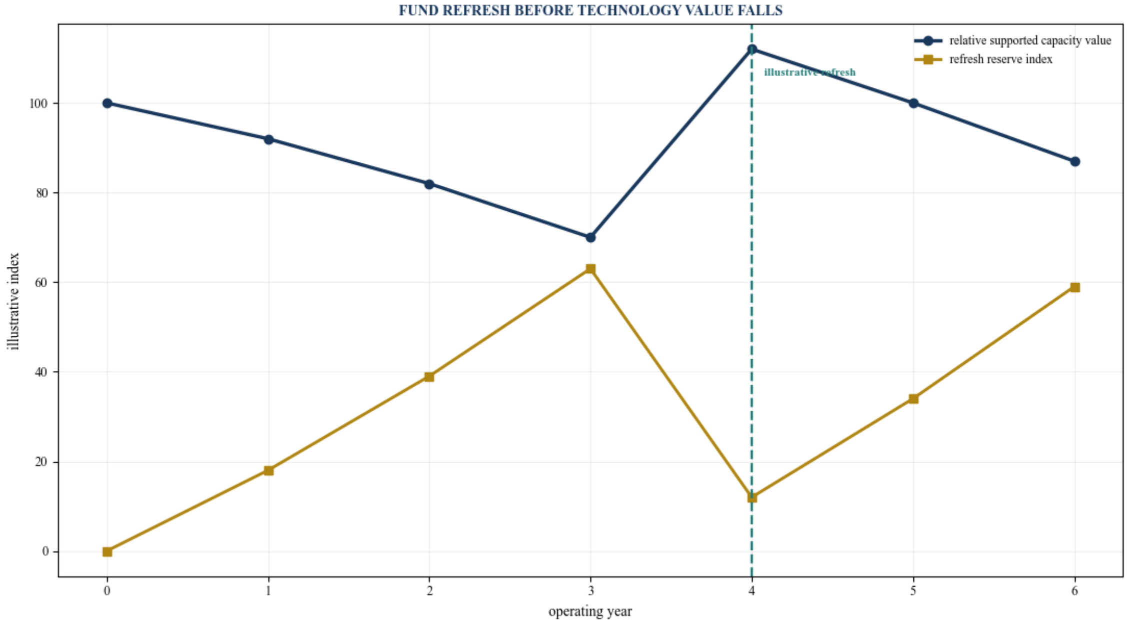

Author framework. Dates and percentages are illustrative and require project-specific approval.

13. Structure contracts around service

Construction completion is only one milestone. Service acceptance should cover power, cooling, cluster performance, network, storage, software, cyber security, scheduler, metering, workload completion and recovery.

The contract package can include an engineering and construction agreement, technology supply, software licences, managed service, power and network arrangements, anchor capacity, government service agreement and user terms.

Interface responsibility needs one accountable integrator or a detailed matrix. Facilities, hardware, networking, software and security vendors can each meet their narrow specification while the end-to-end workload fails.

Performance remedies should relate to affected service. Relevant measures include available accelerator hours, job completion, queue delay, storage throughput, security events, recovery time and loss of reserved capacity. A cap on liability should be assessed against foreseeable service and replacement exposure.

14. Allocate sovereign and supply-chain risk

Sovereign control can be weakened by extraterritorial law, remote administration, proprietary control planes, licence termination, export controls, sanctions, foreign ownership, single-source support and inaccessible cryptographic keys.

The diligence file should map every critical supplier, jurisdiction, subcontractor, licence, remote access path, update mechanism, data flow and termination right. It should identify which functions can continue when a supplier, country or network becomes unavailable.

Mitigations include local operating authority, source escrow where feasible, documented interfaces, multiple network paths, alternative suppliers, spares, portable software, offline recovery, transition assistance and approved substitution rights. Some risks cannot be fully transferred and should be accepted explicitly by the sovereign sponsor.

Security should cover physical assets, firmware, supply chain, identity, privileged access, data, models, training artefacts, monitoring and incident response. Security obligations should remain compatible with user access and operational support.

Table 3. Sovereign compute diligence file

Workstream	Minimum evidence	Investment gate
Sovereign control	ownership, governance, data, keys, management plane, remote access and exit map	required decision rights are legally and operationally enforceable
Demand	benchmarked workloads, user budgets, allocations and anchor terms	phased capacity supported by evidenced use
Site and utilities	title, planning, power, cooling, water, network and resilience	service envelope achievable in normal and degraded states
Technology	benchmarks, interoperability, supply, support, export and refresh	delivered portfolio meets workload and lifecycle requirements
Service layer	identity, scheduler, storage, software, cyber, support and billing	end-to-end user service accepted and measurable
Finance	capital, payment authority, revenue, subsidy, reserves, debt and downside	dependable cash covers operation, debt and refresh
Public benefit	outcome baseline, attribution, reporting and evaluation	benefits have accountable owners and measurable indicators

Each workstream should have an accountable owner and current evidence before approval.

15. Govern data, models and intellectual property

Compute access can expose sensitive data, model weights, prompts, source code, research results and commercial intellectual property. User terms should define ownership, permitted use, confidentiality, retention, deletion, export and incident response.

The national operator should not assume rights to user artefacts beyond those expressly required for operation, security and reporting. Publicly funded projects may carry separate open-access or national-benefit conditions.

Training data and models need provenance and approval suitable for their use. Protected workloads may require isolated environments, controlled administrators and restricted support. Shared systems require strong tenant separation and monitoring.

Intellectual-property conditions should support the programme objective. Excessive claims over user IP can suppress participation. Weak benefit conditions can allow subsidised capacity to create limited national value. The balance should be established transparently in programme rules and contracts.

16. Measure utilisation and useful output

Accelerator utilisation is necessary but incomplete. A system can show high technical utilisation while running low-value jobs, repeated failures or workloads that would be cheaper elsewhere.

The dashboard should combine installed capacity, available hours, allocated hours, started jobs, completed jobs, queue time, failed jobs, energy, cost, user concentration and outputs. Output measures can include trained models, published research, deployed products, public-service use, patents, revenue, users trained and follow-on investment, with appropriate attribution limits.

Capacity reserved for national emergencies or protected work should be reported separately. Low use may be intentional insurance, though its cost and purpose should remain visible.

NAIRR's published two-year update reported more than 600 research teams and 6,000 students supported across the United States.[12] Such participation measures demonstrate reach. Investment evaluation also needs cost, completed workloads and outcome evidence.

17. Test a hypothetical financing case

The following case demonstrates the method. Every number is a hypothetical management assumption. It does not describe an existing national programme or expected market return.

Assume a project company develops a 24 MW initial compute campus with a public service layer. Total initial capital is USD 690 million. The sponsor provides land and enabling works, a public agency commits an availability payment for protected and research capacity, and commercial anchor users reserve capacity. Senior debt funds part of long-life facility and contracted service assets. Technology capital amortises more quickly and receives a refresh reserve.

The hypothetical management assumptions include 55 per cent average useful accelerator utilisation in year one, rising to 72 per cent in year three; 85 per cent service availability for commercial sale after protected and maintenance reservations; and a four-year major accelerator refresh.

Table 4. Hypothetical sovereign compute financing case

Item	Hypothetical management assumption	Underwriting treatment
Initial capacity	24 MW compute campus	phase two requires evidenced contracted demand
Initial capital	USD 690 million	USD 260m facility, USD 360m technology, USD 70m service and contingency
Funding	USD 240m public capital, USD 170m sponsor equity, USD 280m senior debt	debt limited to dependable service cash and conservative collateral
Public service payment	USD 72 million per year for 10 years	subject to authority, appropriation, performance and termination review
Anchor user receipts	USD 38 million in year one, rising to USD 66 million in year three	only signed minimum commitments enter the base case
Other usage revenue	USD 18 million in year one, rising to USD 43 million in year three	scenario-weighted until bookings and payment history emerge
Operating cost	USD 71 million in year one	power, staff, software, support, network, maintenance and security stressed separately
Refresh reserve	USD 34 million in year one, indexed thereafter	restricted cash for technology replacement and migration
Debt service	USD 43 million in year one	tested after operating cost and refresh reserve

Every value is a hypothetical management assumption created solely to demonstrate the method.

Under these assumptions, year-one cash receipts are USD 128 million and cash operating cost plus refresh reserve is USD 105 million, leaving USD 23 million before debt service. This does not cover assumed year-one debt service of USD 43 million. The structure would therefore require additional contracted payment, lower debt, capitalised ramp support or a different phasing plan.

By year three, the hypothetical receipts reach USD 181 million. If operating cost is USD 82 million and the refresh reserve is USD 38 million, cash before debt service is USD 61 million. The year-three case covers the assumed USD 45 million debt service 1.36 times. This is an illustrative calculation, not a forecast.

The case shows that technology refresh cannot be treated as optional capital after debt service. A project can appear financeable when refresh is omitted and become dependent on future emergency appropriations.

18. Stress demand, technology and public cash

The downside model should combine risks that can occur together. Weak user adoption can coincide with rapid technology change. Power constraints can delay capacity while debt interest accrues. Export restrictions can increase equipment price or prevent refresh. A public budget change can reduce service payments or subsidies.

Demand stress should test lower allocations, slower utilisation, concentration, credit loss and price decline. Technology stress should test delayed delivery, underperformance, incompatibility, accelerated refresh and low residual value. Infrastructure stress should test power delay, cooling derating and network failure.

The model should also test policy and control failure: changed eligibility, procurement challenge, appropriation delay, foreign restriction, cyber incident and termination of contributed resources.

Table 5. Sovereign compute stress matrix

Stress	Hypothetical management case	Transmission	Required response
Utilisation ramp delay	year-three useful utilisation reaches 56% rather than 72%	lower usage revenue and weaker unit economics	defer expansion, secure anchors and redesign access
Public payment delay	90-day receipt delay	working-capital and covenant pressure	liquidity reserve, payment mechanics and cure period
Power delay	8 MW arrives 12 months late	idle equipment, lost service and interest carry	staged procurement, long-stop rights and alternative capacity
Faster refresh	major refresh after three years rather than four	higher reserve and earlier capital	increase reserve, shorten debt and preserve substitution
Price decline	commercial rate falls 25%	revenue compression despite demand	minimum commitments, differentiated service and cost reset
Supplier restriction	key accelerator or software unavailable	capacity and continuity loss	multi-vendor plan, portable stack and approved alternatives
Cyber isolation	service segment unavailable for 30 days	lost revenue, recovery cost and confidence	segmented architecture, tested recovery and incident liquidity

Every scenario is a hypothetical management assumption and should be replaced with project evidence.

19. Design the capital stack

Public capital is suited to non-commercial capacity, public benefit, early enabling works and risks the state chooses to retain. Sponsor equity should absorb construction, integration, utilisation and operating risks appropriate to its control.

Senior debt should attach to dependable cash, durable assets and enforceable contracts. Tenor should reflect the shorter of payment visibility, asset life and refresh risk. Technology leases or vendor finance can match equipment life but need careful treatment of control, replacement, default and data.

Availability payments can support project finance when service and deduction mechanics are objective, payment authority is clear and termination compensation is bankable. User take-or-pay contracts can support commercial capacity when workload, credit, price, power pass-through and migration are defined.

Grants and subsidies should be transparent. Their conditions, repayment triggers, stacking rules, audit and outcome requirements can affect cash and control. Canada's access and infrastructure programmes illustrate the need to distinguish stated programme funding from an executed legal obligation.[9][10]

20. Build governance around decisions

The programme needs accountable owners for policy, investment, security, technology, allocation, operation and evaluation. A single board can oversee the platform while delegating protected decisions to authorised committees.

Reserved matters can include control changes, foreign suppliers, material data arrangements, technology substitution, capacity allocation, pricing, debt, refresh, related-party contracts, security exceptions and termination.

Independent technical assurance should test design, benchmark evidence, integration, cyber security, service acceptance and refresh plans. Financial assurance should verify costs, metering, subsidy, reserves and performance payments. Users need a fair appeals and incident process.

Governance should preserve speed. Pre-approved architecture patterns, delegated allocation bands and objective performance standards can reduce recurring approvals while maintaining control.

21. Execute through 180 days

Development should progress from sovereign outcome and demand evidence to delivery structure, market engagement, integrated design, finance and final investment decision.

Figure 6. Sovereign compute investment gate

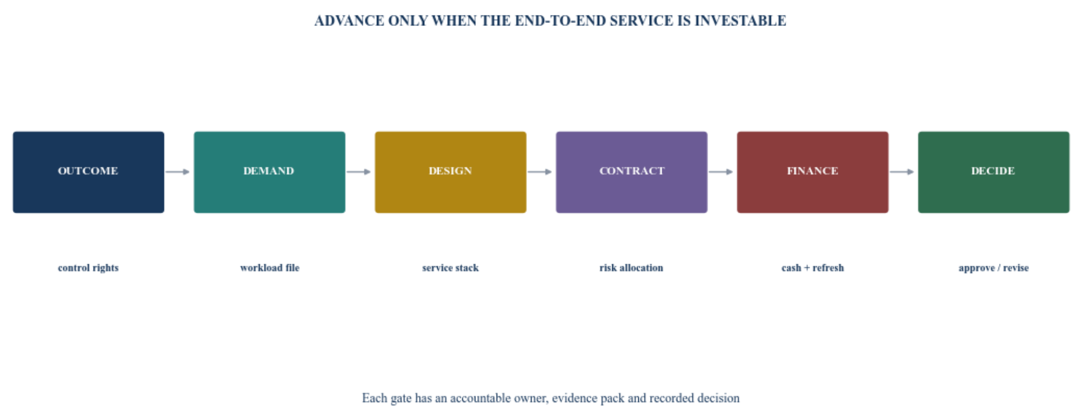

Author framework. The programme advances when control, service, demand, technology and finance evidence are coherent.

Table 6. 180-day development sequence

Days	Workstream	Principal output	Gate
1 to 20	sovereign outcome	use cases, control perimeter, authorities and success measures	mandate approved
21 to 45	demand discovery	segmented users, benchmarks, budgets and security classes	capacity range evidenced
46 to 70	delivery options	public, concession, joint-venture, reserved and federated cases	preferred structure selected
71 to 95	integrated design	site, power, cooling, network, technology and service layer	reference architecture accepted
96 to 120	market engagement	supplier evidence, availability, benchmarks, prices and terms	executable market response obtained
121 to 145	contracts and governance	control, interfaces, access, performance, security, refresh and exit	risk allocation accepted
146 to 165	finance and evaluation	capital stack, payments, reserves, downside and public benefit	financeable case established
166 to 180	investment decision	final evidence, conditions, implementation plan and dashboard	approve, revise or decline

Sequencing is indicative and should be adapted to procurement authority, site readiness and market response.

22. Set the investment decision

The committee should approve the sovereign outcome, control perimeter, demand envelope, delivery model, site, power, cooling, network, accelerator portfolio, service layer, access rules, security, contracts, capital stack, refresh plan and evaluation framework.

Conditions should cover legal authority, appropriations, procurement, planning, utility capacity, export and supplier permissions, benchmark acceptance, security accreditation, anchor commitments, integrated service testing, reserves, insurance and transition rights.

The programme should pause or resize when demand remains unbenchmarked, power is unavailable, control depends on unenforceable assumptions, software is not portable, refresh lacks funding, service interfaces lack accountability or debt depends on speculative usage.

A phased approval can preserve option value. The first phase should establish useful service, demand evidence and operating capability. Expansion should follow contracted or approved workloads, delivered power, measured performance and a refreshed technology case.

23. Limitations and conclusion

Artificial-intelligence hardware, software, networking, workload demand, export controls, public policy, cyber threats, energy systems, environmental constraints, procurement and capital markets can change quickly. Investment decisions require current evidence from authorised public bodies, users, utilities, technology providers, operators, security specialists, engineers, lawyers, insurers, lenders and qualified advisers.

The cited Canadian, European Union, United Kingdom, Indian, United States and Singapore sources describe programmes within their own mandates. They do not establish the legal authority, cost, capacity, performance, demand or return of another project.

Every numerical value and result in the worked case is a hypothetical management assumption. No actual sovereign compute project, public payment, contract, financing, utilisation rate or investment result is claimed.

Sovereign compute becomes an investable asset when national control, user service and financial durability reinforce one another. Physical location is one part of the control perimeter. Productive capacity requires the complete chain from electricity and heat rejection through accelerators, data, software, security, allocation and user support.

The investment case should be built from benchmarked workloads, enforceable control, integrated service acceptance, dependable payments, conservative residual value and a funded refresh cycle. Public benefits should be measured through accountable outcomes. Commercial cash should remain distinct and traceable to contracts or authorised payments.

References

[1] Innovation, Science and Economic Development Canada, Program Guide: Artificial Intelligence Sovereign Compute Infrastructure Program, 2026. https://ised-isde.canada.ca/site/ised/en/program-guide-artificial-intelligence-sovereign-compute-infrastructure-program-scip

[2] European High Performance Computing Joint Undertaking, AI Factories Services, official programme page accessed 13 August 2026. https://www.eurohpc-ju.europa.eu/ai-factories/ai-factories-services_en

[3] European High Performance Computing Joint Undertaking, The EuroHPC JU Selects Additional AI Factories to Strengthen Europe's AI Leadership, 12 March 2025. https://www.eurohpc-ju.europa.eu/eurohpc-ju-selects-additional-ai-factories-strengthen-europes-ai-leadership-2025-03-12_en

[4] United States National Science Foundation, National Artificial Intelligence Research Resource Pilot, 8 March 2024. https://www.nsf.gov/geo/updates/national-artificial-intelligence-research-resource-nairr

[5] United Kingdom Department for Science, Innovation and Technology and UK Research and Innovation, AI Research Resource, updated 17 July 2026. https://www.gov.uk/government/publications/ai-research-resource

[6] IndiaAI, IndiaAI Compute Portal, official portal accessed 13 August 2026. https://compute.indiaai.gov.in/login

[7] IndiaAI, Compute Price Calculator, official portal accessed 13 August 2026. https://compute.indiaai.gov.in/indiaaipricecalculator

[8] Innovation, Science and Economic Development Canada, Canada to Drive Billions in Investments to Build Domestic AI Compute Capacity at Home, 5 December 2024. https://www.canada.ca/en/innovation-science-economic-development/news/2024/12/canada-to-drive-billions-in-investments-to-build-domestic-ai-compute-capacity-at-home.html

[9] Innovation, Science and Economic Development Canada, AI Sovereign Compute Infrastructure Program, updated 1 June 2026. https://ised-isde.canada.ca/site/ised/en/ai-sovereign-compute-infrastructure-program

[10] Innovation, Science and Economic Development Canada, Program Guide: AI Compute Access Fund, updated 24 June 2025. https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy/ai-compute-access-fund/program-guide-ai-compute-access-fund

[11] United Kingdom Department for Science, Innovation and Technology, UK Compute Roadmap, 2026. https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap

[12] United States National Science Foundation, NAIRR at 2 Years: Advancing American Artificial Intelligence Innovation and Leadership, 19 March 2026. https://www.nsf.gov/cise/updates/nairr-2-years-advancing-american-artificial-intelligence

[13] Singapore Economic Development Board and Infocomm Media Development Authority, Launch of Second Data Centre Call for Application, 1 December 2025. https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/factsheets/2025/launch-of-second-data-centre

About the Author

Chennakeshav Adya is an independent researcher and Managing Partner of Matchpoint Partners. His research focuses on investment strategy, capital formation, transaction execution, governance and operating-model design across international markets.