# NORAI

North AI: Canada's Sovereign Platform for
AI-Accelerated Scientific Discovery

Semanttica Web Technologies (SWT) Inc.

Version 2.0 – December 2025

## Contents

## Executive Summary

The United States launched the Genesis Mission on 25 November 2025 to double the productivity of American science within a decade using closed-loop AI on sovereign data and supercomputers. Canada cannot afford to watch from the sidelines.

**NORAI** is the Made-in-Canada answer: a fully sovereign, bilingual, Indigenous-consent-aware AI-for-science platform that integrates Canadian scientific datasets, laboratories, and supercomputers from coast to coast to coast under one secure national fabric.

**Canada doesn't need to copy Genesis. Canada needs to build something Genesis cannot replicate:** a platform with Indigenous data governance (OCAP), bilingual by design, privacy-preserving federation across provinces, and exportable to allied nations.

Built and operated in Canada by Canadian companies, NORAI will:

- **Create 1,000+ direct jobs** and 3,000+ indirect jobs over 5 years
- **Redirect $500M+/year** currently flowing to US cloud providers back into Canadian infrastructure
- **Retain AI talent** by providing compelling domestic opportunities
- **Enable $33B+ in economic value** across critical minerals, health, climate, and energy sectors (23x ROI)

We are not asking for permission to start - we are asking for partnership to mobilize Canadian talent and build sovereign AI infrastructure.

# 1. The Strategic Imperative

## 1.1 The Global AI-for-Science Race

| Nation | Initiative | Compute Commitment | Investment |
|---|---|---|---|
| United States | DOE Genesis Mission | 200 MW dedicated | $3B+ |
| China | New Generation AI 2030 | 3 closed-loop platforms | Est. $10B+ |
| European Union | EuroHPC + AI4Science | 5 exascale systems | EUR 8B |
| **Canada** | **Fragmented efforts** | No integrated platform | Fragmented |

Table 1: Global AI-for-Science Initiatives

## 1.2 The Cost of Inaction

- Canadian scientists waste **40-60%** of their time on data wrangling instead of discovery
- Research talent migration to US institutions accelerating post-Genesis announcement
- **$500M+/year** flows to US cloud providers for AI workloads
- Critical sectors (climate, health, energy) lack integrated AI infrastructure
- Indigenous data sovereignty concerns remain unaddressed

## 1.3 Canada's Unique Opportunity

Canada doesn't need to match Genesis dollar-for-dollar. Canada needs to build capabilities Genesis **cannot replicate**:

| Capability | Genesis (US) | NORAI (Canada) |
|---|---|---|
| Indigenous data governance | None | OCAP-native (world first) |
| Bilingual platform | No | EN/FR/Inuktitut-ready |
| Provincial federation | N/A | Privacy-preserving across 13 jurisdictions |
| Export to allies | Unlikely | Five Eyes, EU, Francophone nations |
| Open science ethos | DOE-controlled | FAIR principles, open APIs |

**NORAI closes the gap by building what only Canada can build.**

# 2. Made-in-Canada Value Proposition

## 2.1 Economic Impact

| Metric | Year 1 | Year 3 | Year 5 |
|---|---|---|---|
| Direct jobs | 180 | 480 | 1,000+ |
| Indirect jobs | 540 | 1,440 | 3,000+ |
| Canadian company contracts | $20M | $100M | $300M |
| Compute value unlocked | $50M | $150M | $300M |
| Research productivity gain | 10% | 30% | 50% |

## 2.2 Keeping Value in Canada

| Current State | With NORAI |
|---|---|
| $500M+/year to US cloud providers | Redirected to Canadian infrastructure |
| AI talent emigrating to US | Compelling domestic opportunities |
| Research data exported for analysis | Processed on sovereign infrastructure |
| Foreign platforms control access | Canadian-governed, Canadian-controlled |

## 2.3 Canadian Company Opportunities

NORAI creates a platform for Canadian technology companies to compete globally:

| Sector | Canadian Companies to Engage |
|---|---|
| **Cloud/Compute** | Shopify, OVHcloud (Montreal), IBM Canada, TELUS, Bell |
| **AI/ML** | Cohere (Toronto), Sanctuary AI (Vancouver), Untether AI |
| **Cybersecurity** | BlackBerry/Cylance, Arctic Wolf, eSentire |
| **Quantum** | Xanadu (Toronto), D-Wave (Burnaby), Photonic Inc |
| **Data/Analytics** | ThinkData Works, Daisy Intelligence, Coveo |
| **Health Tech** | WELL Health, Think Research, MedStack |
| **Indigenous Tech** | First Nations Technology Council, Animikii |
| **Geospatial** | MDA, Ecopia AI, Geotab |

## 2.4 Talent Retention and Development

| Initiative | Impact |
|---|---|
| Graduate research positions | 200+ funded positions across Canadian universities |
| Industry partnerships | Co-op and internship programs with Canadian companies |
| Indigenous technology training | Dedicated programs in partnership with Indigenous institutions |
| Regional development | Positions distributed across provinces, not just Toronto/Montreal |

# 3. Platform Architecture

NORAI implements a five-layer sovereign architecture designed for scientific discovery at national scale.

## 3.1 Layer 1: National Semantic Layer

**Purpose:** Unified knowledge representation across all Canadian scientific domains

| Component | Specification |
|---|---|
| Knowledge Graph | 50M+ scientific entities (target) |
| Ontologies | Aligned with FAIR principles, W3C standards |
| Languages | English, French, Inuktitut-ready |
| Entity Types | Researchers, publications, datasets, instruments, samples, experiments |
| Integration | CrossRef, ORCID, DataCite, Canadian institutional repositories |

## 3.2 Layer 2: Federated Data Lakehouse

**Purpose:** Distributed sovereign data storage with automatic residency enforcement

| Component | Specification |
|---|---|
| Architecture | Delta Lake on object storage (S3-compatible) |
| Geographic Zones | Montreal, Toronto, Calgary, Vancouver, Yellowknife |
| Compliance | PIPEDA, provincial health data regulations, OCAP |
| Capacity | Petabyte-scale per zone, exabyte aggregate |
| Data Formats | Parquet, Zarr, HDF5, DICOM, FASTA, netCDF |

## 3.3 Layer 3: Secure Multi-Party Compute Enclave

**Purpose:** Privacy-preserving computation across institutional and jurisdictional boundaries

## 3.4 Layer 4: Borealis Model Family

**Purpose:** Science-specialized foundation models trained on Canadian data

## 3.5 Layer 5: NORAI Studio

**Purpose:** End-to-end research automation interface

| Component | Specification |
|---|---|
| MPC Framework | OpenMPC with Canadian security extensions |
| Network Security | SCION-based path-aware networking |
| Access Control | Consent-as-code governance tokens |
| Differential Privacy | Configurable ε (epsilon) budgets per query |
| Audit | Immutable computation provenance logs |

| Model | Parameters | Training Data | Capabilities |
|---|---|---|---|
| Borealis-8B | 8 billion | Canadian scientific literature, patents | Text understanding, code generation |
| Borealis-70B | 70 billion | + multi-modal (images, spectra, time series) | Scientific reasoning, hypothesis generation |
| Borealis-405B | 405 billion | + closed-loop experimental data | Autonomous experiment design |

# 4. Year 1 Prototype Priorities

## 4.1 Computome - The National Compute Gateway

**Problem:** Alliance clusters at 60

## 4.2 Federated Data Catalog

**Problem:** Researchers spend 40-60

## 4.3 Consent-as-Code (OCAP Platform)

**Problem:** No platform automates Indigenous data governance at scale

## 4.4 Provincial Data Sharing Protocol

**Problem:** Cross-provincial research requires 6+ months of legal paperwork

## 4.5 Critical Minerals AI

**Problem:** Manual geological analysis; supply chain dependence on foreign sources

## 4.6 PermafrostGPT

**Problem:** $8-12B in northern infrastructure at risk from permafrost degradation

## 4.7 Privacy-Preserving Health Analytics

**Problem:** Health research requires centralizing sensitive data (provinces refuse)

## 4.8 NORAI Data Gateway - Unified Open Science API

**Problem:** 30+ Canadian open science platforms exist but are siloed; researchers waste 40-60

| Component | Specification |
|---|---|
| Interface | Natural language to experiment design |
| Automation | Robotic laboratory execution protocols |
| Provenance | Full experiment lineage tracking |
| Integration | TRIUMF, Canadian Light Source, Ocean Networks Canada |

| **Aspect** | **Specification** |
|---|---|
| Solution | Unified gateway + intelligent dispatcher for all Alliance resources |
| Value | $50M/year in unlocked compute; keeps AI training in Canada |
| Stakeholders | Digital Research Alliance, ISED, universities, startups |
| Deliverables | Dashboard, job dispatcher, queue optimization |
| Canadian jobs | 10-15 software engineers |

| **Aspect** | **Specification** |
|---|---|
| Solution | Search API federating metadata from StatsCan, NRCan, ECCC, CIHR |
| Value | Accelerates every Canadian researcher |
| Stakeholders | All federal science departments |
| Deliverables | Search API, web interface, FAIR metadata harvesting |
| Canadian jobs | 15-20 developers, data engineers |

| **Aspect** | **Specification** |
|---|---|
| Solution | Machine-readable consent policies with community control |
| Value | World's first - competitive advantage vs. Genesis |
| Stakeholders | CIRNAC, ITK, AFN, FNIGC |
| Deliverables | Consent portal, policy engine, audit system |
| Canadian jobs | Indigenous tech specialists, governance experts |

| **Aspect** | **Specification** |
|---|---|
| Solution | Automated compliance for PIPEDA + provincial health acts |
| Value | Enables pan-Canadian health research |
| Stakeholders | Provincial health ministries, CIHR |
| Deliverables | Compliance automation, legal templates, workflow system |
| Canadian jobs | Legal tech, privacy engineers |

| **Aspect** | **Specification** |
|---|---|
| Solution | AI analysis of satellite + geophysics for deposit identification |
| Value | $15B+ export potential |
| Stakeholders | NRCan, mining sector, Indigenous communities |
| Deliverables | Analysis pipeline, probability maps, prospect database |
| Canadian jobs | Geoscientists, ML engineers, remote sensing |

| **Aspect** | **Specification** |
|---|---|
| Solution | 1-km resolution 10-year forecasts using GSC + satellite data |
| Value | Infrastructure protection, Arctic sovereignty |
| Stakeholders | NRCan, territorial governments, Transport Canada |
| Deliverables | Forecast system, risk maps, API |
| Canadian jobs | Climate scientists, infrastructure engineers |

| **Aspect** | **Specification** |
|---|---|
| Solution | Federated learning with differential privacy |
| Value | 10M+ patient records analyzable without data movement |
| Stakeholders | Provincial health, BC Cancer, Ontario Health |
| Deliverables | Federated learning framework, privacy guarantees |
| Canadian jobs | Privacy engineers, health informatics |

| **Aspect** | **Specification** |
|---|---|
| Solution | Unified API wrapper federating all Canadian open science platforms with native ML/AI integration |
| Value | Single interface to 30+ platforms; direct PyTorch/TensorFlow data loading; automatic format conversion |
| Stakeholders | CONP, Ocean Networks Canada, CADC, Polar Data Catalogue, FRDR, all open science platforms |
| Deliverables | Python SDK, REST API, platform adapters, data propagation bus |
| Canadian jobs | 20-25 software engineers, data engineers, DevOps |
| Investment | $8M Year 1 |

### 4.9 Compute Queue Predictor & Optimizer (Quick Win)

**Problem:** Researchers submit jobs blindly, wait days, jobs fail. Alliance RAC 2024 fulfilled only 43

| Aspect | Specification |
|---|---|
| Solution | ML model predicting wait times across clusters, suggesting optimal submission times and resources, auto-retry for failed jobs |
| Value | Better utilization of existing $225M infrastructure; reduced researcher frustration; fewer wasted compute hours |
| Stakeholders | Digital Research Alliance, all HPC users (10,000+ researchers) |
| Deliverables | Wait time predictor API, CLI tool, Slurm plugin, dashboard |
| Canadian jobs | 5-8 ML engineers, backend developers |
| Investment | $1M Year 1 |
| Timeline | 3 months to MVP |

### 4.10 BorealisChat - Bilingual Scientific Assistant (High Visibility)

**Problem:** No Canadian-specific, bilingual AI assistant for researchers. All use US-based tools (ChatGPT, Claude) with no understanding of Canadian context, Tri-Council requirements, or French scientific literature.

| Aspect | Specification |
|---|---|
| Solution | RAG-based chatbot trained on Canadian scientific literature, government reports, Tri-Council funding guidelines, bilingual (EN/FR) |
| Value | High visibility "Made in Canada" AI; serves underserved French-speaking research community; demonstrates sovereignty |
| Stakeholders | All Canadian researchers, ISED, Canadian Heritage (bilingualism mandate) |
| Deliverables | Web interface, API, Slack/Teams integration, mobile app |
| Canadian jobs | 10-15 ML engineers, NLP specialists, bilingual content specialists |
| Investment | $4M Year 1 |
| Timeline | 6 months to public beta |

### 4.11 Carbon-Aware Compute Scheduler (Climate Win)

**Problem:** Alliance clusters at 80

| Aspect | Specification |
|---|---|
| Solution | Route batch jobs to lowest-carbon regions; time-shift non-urgent jobs to low-carbon periods; carbon footprint reporting |
| Value | 20-30% carbon reduction; aligns with federal climate policy; ESG reporting for institutions |
| Stakeholders | ECCC, Digital Research Alliance, universities, federal science departments |
| Deliverables | Scheduler plugin, carbon API integration, reporting dashboard |
| Canadian jobs | 5-7 systems engineers, sustainability specialists |
| Investment | $1.5M Year 1 |
| Timeline | 4 months to production |

### 4.12 Grant Writing Assistant (High Demand)

**Problem:** Researchers spend weeks writing NSERC/CIHR/SSHRC grants. 70

### 4.13 Research Impact Tracker

**Problem:** Canadian researchers can't easily demonstrate impact of publicly-funded research. Tri-Council needs better metrics to justify $4B+ annual investment.

| Aspect | Specification |
|---|---|
| Solution | AI assistant understanding Tri-Council formats, trained on successful grants (anonymized), reviewer criteria, common pitfalls |
| Value | Democratizes grant access; saves 40+ hours per application; improves success rates for underserved institutions |
| Stakeholders | Tri-Council (NSERC, CIHR, SSHRC), universities, early-career researchers |
| Deliverables | Web application, document analysis, suggestion engine, compliance checker |
| Canadian jobs | 8-12 ML engineers, UX designers, grant specialists |
| Investment | $2.5M Year 1 |
| Timeline | 4 months to pilot |

| Aspect | Specification |
|---|---|
| Solution | Automated tracking of dataset citations, code reuse, downstream discoveries, policy citations, media mentions |
| Value | Justifies public investment; helps researchers in grant renewals; identifies high-impact research for funding |
| Stakeholders | Tri-Council, universities, research offices, government |
| Deliverables | Impact dashboard, researcher profiles, institutional reports, DOI/ORCID integration |
| Canadian jobs | 6-10 data engineers, bibliometrics specialists |
| Investment | $2M Year 1 |
| Timeline | 3 months to MVP |

### 4.14 Cross-Institutional Collaboration Matcher

**Problem:** Researchers don't know who else in Canada works on similar problems. Duplicate efforts waste resources. Grant applications weaker without collaborators.

| Aspect | Specification |
|---|---|
| Solution | Knowledge graph of researchers, publications, datasets, grants with AI-powered similarity matching and collaboration recommendations |
| Value | Enables new collaborations; reduces duplication; strengthens grant applications; builds Canadian research networks |
| Stakeholders | Tri-Council, universities, research networks, early-career researchers |
| Deliverables | Researcher discovery portal, collaboration recommendations, research landscape maps |
| Canadian jobs | 5-8 knowledge graph engineers, NLP specialists |
| Investment | $1.5M Year 1 |
| Timeline | 3 months to beta |

### 4.15 World-Leading Innovations (Priorities 15-21)

The following prototypes position Canada as a global leader in AI for science by building capabilities that no other nation is currently developing at scale.

### 4.16 Arctic Digital Twin (Borealis-Earth) - WORLD FIRST

**Problem:** Canada has 40

| Aspect | Specification |
|---|---|
| Solution | First km-resolution digital twin of Canadian Arctic and boreal ecosystems, integrating permafrost, sea ice, wildlife corridors, and infrastructure |
| Value | Enables predictive infrastructure planning, supports Arctic sovereignty, powers climate adaptation policy |
| Stakeholders | ECCC, NRCan, DND/CAF, territorial governments, Inuit Tapiriit Kanatami |
| Deliverables | Interactive 3D visualization, API for simulations, scenario planning tools, policy decision support |
| Canadian jobs | 25-35 climate scientists, visualization engineers, Arctic specialists |
| Investment | $12M Year 1 |
| Timeline | 12 months to operational prototype |
| Competitive Position | No equivalent exists globally; positions Canada as Arctic AI leader |

### 4.17 Autonomous Materials Discovery Lab (Self-Driving Science) - LEAPFROG

**Problem:** US DOE Genesis Mission is building autonomous labs at Argonne/Berkeley. Canada has no equivalent. Manual materials discovery takes 10-20 years per new material. Critical minerals supply chain depends on accelerating discovery.

| Aspect | Specification |
|---|---|
| Solution | Closed-loop AI + robotics lab for critical minerals and clean energy materials - AI designs experiments, robots execute, ML analyzes, loop repeats 24/7 |
| Value | 10x faster materials discovery; reduces critical minerals import dependence; positions Canada in global battery/clean energy race |
| Stakeholders | NRC, NRCan, Canadian Light Source, TRIUMF, mining sector |
| Deliverables | Physical autonomous lab facility, AI experiment planner, robotic integration, materials database |
| Canadian jobs | 30-40 roboticists, materials scientists, ML engineers, lab technicians |
| Investment | $15M Year 1 (facility + equipment + team) |
| Timeline | 12 months to first autonomous discovery cycle |
| Competitive Position | Matches Genesis Mission capability; first in Canada |

### 4.18 Indigenous AI Governance Platform (CARE/OCAP Automation) - WORLD FIRST

**Problem:** No platform anywhere automates Indigenous data governance at scale. CARE Principles and OCAP exist as frameworks but require manual implementation. Indigenous communities lack technical tools to enforce data sovereignty.

| Aspect | Specification |
|---|---|
| Solution | Complete technology stack for Indigenous data sovereignty: community-controlled consent portals, automated CARE/OCAP compliance, benefit-sharing smart contracts, data repatriation tools |
| Value | World's first - exportable to Indigenous communities globally; positions Canada as ethical AI leader; addresses reconciliation commitments |
| Stakeholders | ITK, AFN, MNC, FNIGC, CIRNAC, Indigenous tech companies (Animikii, FNTC) |
| Deliverables | Community governance portal, consent workflow engine, benefit-sharing tracker, audit dashboard, training program |
| Canadian jobs | 15-20 Indigenous tech specialists, governance experts, community liaisons |
| Investment | $5M Year 1 |
| Timeline | 8 months to pilot with 3 communities |
| Competitive Position | Absolute world first; no equivalent exists; exportable IP |

### 4.19 Canadian Scientific Foundation Model (Borealis-Science) - SOVEREIGN AI

**Problem:** All scientific AI assistants (GPT-4, Claude, Gemini) are US-based, trained on global data, with no understanding of Canadian context, French scientific literature, or Indigenous knowledge. Canada has no sovereign scientific AI.

| Aspect | Specification |
|---|---|
| Solution | Domain-specific 70B+ parameter foundation model trained exclusively on Canadian scientific corpus, bilingual (EN/FR), with optional Indigenous language support |
| Value | Sovereign AI capability; serves underserved French research community; demonstrates "Made in Canada" AI; can be exported to Francophone nations |
| Stakeholders | ISED, NRC, universities, Cohere (potential partner), Canadian Heritage |
| Deliverables | Pre-trained foundation model, fine-tuning toolkit, API access, integration with NORAI platform |
| Canadian jobs | 40-50 ML researchers, data engineers, linguists, compute infrastructure |
| Investment | $20M Year 1 (compute-intensive) |
| Timeline | 12 months to Borealis-Science 70B release |
| Competitive Position | First national sovereign scientific foundation model; differentiator vs Genesis |

### 4.20 Federated Genomics Network - PAN-CANADIAN HEALTH

**Problem:** Canada's provincial health system creates data silos. Cancer/genomics research requires combining data across provinces, but privacy laws prevent centralization. No existing platform enables federated genomic analysis at national scale.

| Aspect | Specification |
|---|---|
| Solution | Privacy-preserving federated learning network connecting all provincial cancer registries and genomic databases - analyze 10M+ patient records without any data leaving provincial jurisdiction |
| Value | Enables pan-Canadian precision medicine research impossible today; 3+ new therapeutic targets; model for global federated health AI |
| Stakeholders | CIHR, provincial cancer agencies, Genome Canada, Health Canada |
| Deliverables | Federated learning framework, secure enclaves in each province, genomic analysis pipelines, researcher portal |
| Canadian jobs | 20-25 privacy engineers, health informaticists, genomics specialists |
| Investment | $8M Year 1 |
| Timeline | 10 months to operational network (3 provinces) |
| Competitive Position | Largest federated genomics network in the world; leverages Canada's universal healthcare |

### 4.21 Borealis-Scholar - AI Research Agent - OPEN ALTERNATIVE

**Problem:** FutureHouse and Elicit (US) dominate AI-powered literature synthesis. No open-source, bilingual alternative exists. Canadian researchers use US tools with no Canadian context or French support.

### 4.22 Quantum-Classical Science Bridge - QUANTUM ADVANTAGE

**Problem:** Canada has world-leading quantum companies (Xanadu, D-Wave) but no integration with national HPC. Quantum advantage for science requires hybrid classical-quantum workflows. No national platform exists to make this accessible.

| **Aspect** | **Specification** |
|---|---|
| Solution | Open-source AI research agent with superhuman literature search and synthesis, trained on Canadian research corpus, fully bilingual, understanding of Tri-Council ecosystem |
| Value | Reduces literature review from weeks to hours; serves French research community; open-source builds Canadian AI ecosystem |
| Stakeholders | All Canadian researchers, CARL (Canadian Association of Research Libraries), Érudit |
| Deliverables | Web application, API, VS Code extension, Zotero integration, open-source codebase |
| Canadian jobs | 12-18 ML engineers, NLP specialists, research librarians |
| Investment | $4M Year 1 |
| Timeline | 6 months to public beta |
| Competitive Position | Open-source alternative to US commercial tools; only bilingual option |

| **Aspect** | **Specification** |
|---|---|
| Solution | First national platform seamlessly connecting quantum processors (Xanadu photonic, D-Wave annealing) with Alliance HPC clusters for hybrid scientific workflows |
| Value | Democratizes quantum access; accelerates quantum advantage timeline; keeps quantum talent in Canada; first-mover in quantum-HPC integration |
| Stakeholders | Xanadu, D-Wave, Digital Research Alliance, NRC, universities with quantum programs |
| Deliverables | Unified API, job scheduler for hybrid workflows, quantum simulation fallback, educational resources |
| Canadian jobs | 15-20 quantum software engineers, HPC specialists |
| Investment | $10M Year 1 |
| Timeline | 12 months to hybrid workflow demonstrations |
| Competitive Position | First national quantum-HPC integration; leverages Canadian quantum leadership |

## 4.23 Year 1 Prototype Summary

## 4.24 Core Infrastructure Prototypes (Priorities 1-14)

| Priority | Prototype | Investment | Timeline | Jobs |
|---|---|---|---|---|
| 1 | Computome | $8M | Q1-Q2 | 10-15 |
| 2 | Federated Data Catalog | $6M | Q1-Q3 | 15-20 |
| 3 | Consent-as-Code (OCAP) | $7M | Q2-Q4 | 10-15 |
| 4 | Provincial Data Sharing Protocol | $8M | Q2-Q4 | 12-18 |
| 5 | Critical Minerals AI | $9M | Q2-Q4 | 15-20 |
| 6 | PermafrostGPT | $7M | Q3-Q1 | 10-15 |
| 7 | Privacy-Preserving Health Analytics | $5M | Q3-Q2 | 8-12 |
| 8 | NORAI Data Gateway | $8M | Q1-Q4 | 20-25 |
| 9 | Compute Queue Predictor | $1M | Q1 | 5-8 |
| 10 | BorealisChat | $4M | Q1-Q2 | 10-15 |
| 11 | Carbon-Aware Scheduler | $1.5M | Q1-Q2 | 5-7 |
| 12 | Grant Writing Assistant | $2.5M | Q1-Q2 | 8-12 |
| 13 | Research Impact Tracker | $2M | Q1 | 6-10 |
| 14 | Collaboration Matcher | $1.5M | Q1 | 5-8 |
| **Subtotal** | **14 Prototypes** | **$70.5M** | | **140-200** |

*Quick Win prototypes (9-14) deliver in 3-4 months.*

## 4.25 World-Leading Innovations (Priorities 15-21)

| Priority | Prototype | Investment | Timeline | Jobs | Position |
|---|---|---|---|---|---|
| 15 | Arctic Digital Twin | $12M | Q1-Q4 | 25-35 | **World First** |
| 16 | Autonomous Materials Lab | $15M | Q1-Q4 | 30-40 | Genesis Parity |
| 17 | Indigenous AI Governance | $5M | Q1-Q3 | 15-20 | **World First** |
| 18 | Borealis-Science Foundation Model | $20M | Q1-Q4 | 40-50 | Sovereign AI |
| 19 | Federated Genomics Network | $8M | Q1-Q4 | 20-25 | Global Model |
| 20 | Borealis-Scholar AI Agent | $4M | Q1-Q2 | 12-18 | Open Alternative |
| 21 | Quantum-Classical Bridge | $10M | Q1-Q4 | 15-20 | First National |
| **Subtotal** | **7 World-Leading** | **$74M** | | **157-208** | |

## 4.26 Combined Year 1 Investment

| Category | Prototypes | Investment | Jobs Created |
|---|---|---|---|
| Core Infrastructure | 14 | $70.5M | 140-200 |
| World-Leading Innovations | 7 | $74M | 157-208 |
| **Total Year 1** | **21 Prototypes** | **$144.5M** | **297-408** |

**Strategic Breakdown:**

- **2 World Firsts** - Capabilities no other nation has built
- **1 Genesis Parity** - Matching US autonomous lab capability
- **1 Sovereign AI** - Canadian foundation model independence
- **6 Quick Wins** - Demonstrable value in 3-4 months
- **1 Global Model** - Template exportable to allies

# 5. World-Leading Innovation Technical Specifications

## 5.1 Arctic Digital Twin (Borealis-Earth)

**Objective:** First km-resolution digital twin of the Canadian Arctic and boreal ecosystems

| Aspect | Specification |
|---|---|
| Spatial Resolution | 1 km (Arctic), 250m (infrastructure corridors) |
| Temporal Resolution | Hourly for weather, daily for permafrost, seasonal for ecosystems |
| Coverage | 4.5 million km² (all Canadian Arctic and boreal zones) |
| Forecast Horizon | 10-day weather, 10-year infrastructure risk, 2100 climate scenarios |

**Data Sources:**

| Source | Data Type | Update Frequency |
|---|---|---|
| RADARSAT Constellation | SAR imagery, ice extent | Daily |
| MODIS/VIIRS | Surface temperature, snow cover | Daily |
| GSC Borehole Network | Ground temperature, permafrost depth | Continuous |
| Environment Canada | Weather stations, buoys | Hourly |
| Polar Data Catalogue | Historical records, research data | As available |
| Indigenous Knowledge | Traditional observations (with consent) | Seasonal |

**Model Architecture (Borealis-Earth):**

- **Data Ingestion Layer:**
  - Satellite data (SAR, Optical)
  - Ground Sensors
  - Weather Stations
  - Indigenous Knowledge Portal
- **Physics-Informed ML Core:**
  - Permafrost Model (FourCastNet-based)
  - Sea Ice Model (Graph NN)
  - Ecosystem Model (Species Distribution)
- **Visualization Engine:**
  - 3D Globe (CesiumJS)
  - Risk Maps (Mapbox)
  - Scenario Explorer

**Key Capabilities:**

| Capability | Description |
|---|---|
| Infrastructure Risk Scoring | Real-time risk assessment for roads, pipelines, buildings |
| Shipping Route Optimization | Safe passage planning through Northwest Passage |
| Wildlife Corridor Modeling | Caribou, polar bear migration predictions |
| Community Alerts | Early warning system for hazardous conditions |
| Policy Scenario Testing | "What if" simulations for climate policy decisions |

**Performance Specifications:**

| Metric | Target |
|---|---|
| Simulation speed | 10x real-time (1 year in 36 days) |
| Inference latency | <5 minutes for regional forecast |
| Storage | 2 PB (growing 500 TB/year) |
| Compute | 500 GPU-hours/day continuous operation |

## 5.2 Autonomous Materials Discovery Lab

**Objective:** Self-driving laboratory for critical minerals and clean energy materials discovery

| Aspect | Specification |
|---|---|
| Lab Type | High-throughput materials synthesis and characterization |
| Automation Level | A4-A5 (fully autonomous operation with human oversight) |
| Focus Areas | Battery materials, rare earth alternatives, fusion materials |
| Throughput | 100+ experiments/day (vs. 5/day manual) |

**Physical Infrastructure:**

| Component | Equipment | Purpose |
|---|---|---|
| Synthesis | Robotic powder handling, furnaces, ball mills | Material creation |
| Characterization | XRD, SEM, XRF (automated sample loading) | Property measurement |
| Electrochemistry | Coin cell assembly robot, cycler array | Battery testing |
| Compute | On-site GPU cluster (8x A100) | Real-time ML inference |
| Safety | Inert atmosphere glovebox, fire suppression | Safe operation |

**AI Experiment Planner:** Bayesian optimization with physics-informed priors searches chemical space for promising candidates, filters for safety and synthesizability, and proposes experiment plans with estimated times and safety notes. The system learns from each result and updates the materials database.

**Integration Points:**

## 5.3 Indigenous AI Governance Platform

**Objective:** World's first automated platform for Indigenous data sovereignty (CARE/OCAP compliance)

**Platform Components:**

**Consent Workflow:**

**Benefit Sharing Automation:**

## 5.4 Borealis-Science Foundation Model

**Objective:** Sovereign Canadian scientific foundation model (70B+ parameters)

**Training Data:**

**Model Capabilities:**

**Training Infrastructure:**

**Safety and Alignment:**

## 5.5 Federated Genomics Network

**Objective:** Pan-Canadian federated learning network for genomics/cancer research

| Facility | Integration | Purpose |
|---|---|---|
| Canadian Light Source | Beamline scheduling API | High-resolution characterization |
| TRIUMF | Muon spectroscopy queue | Battery material analysis |
| NRC-ICPET | Computational validation | DFT cross-validation |
| Materials Project | Database sync | Global knowledge sharing |

| Aspect | Specification |
|---|---|
| Frameworks Supported | OCAP (Canada), CARE Principles (global), Māori Data Sovereignty |
| Governance Model | Community-controlled with delegated technical administration |
| Audit | Immutable blockchain-based provenance |
| Languages | English, French, Inuktitut, Cree, Ojibwe (expandable) |

| Component | Technology | Purpose |
|---|---|---|
| Community Portal | Next.js + Keycloak | Self-service governance interface |
| Consent Engine | Smart contracts (Hyperledger Fabric) | Automated policy enforcement |
| Benefit Tracker | PostgreSQL + Grafana | Track research benefits back to community |
| Repatriation Tools | IPFS + encrypted transfer | Secure data return to communities |
| Training Platform | LMS + video conferencing | Capacity building for communities |

| Benefit Type | Tracking Method | Distribution |
|---|---|---|
| Financial | Grant overhead, commercialization royalties | Direct to community fund |
| Knowledge | Publications, presentations | Copies provided, co-authorship offered |
| Capacity | Training, equipment | Tracked, reported annually |
| Acknowledgment | Citations, media | Monitored, alerts to community |

| Aspect | Specification |
|---|---|
| Parameters | 70B (Phase 1), 405B (Phase 2) |
| Architecture | Mixture of Experts (MoE) transformer |
| Languages | English, French (balanced 50/50) |
| Context Window | 128K tokens |
| Training Infrastructure | 100% Canadian (Alliance clusters + OVHcloud) |

| Source | Size | Content |
|---|---|---|
| Canadian Scientific Literature | 2M papers | NRC-CISTI, Érudit, institutional repositories |
| Federal Government Reports | 500K docs | Open Government Portal, departmental publications |
| Patents | 300K | CIPO Canadian patents |
| Tri-Council Guidelines | 10K docs | NSERC, CIHR, SSHRC policies and forms |
| Indigenous Knowledge (with consent) | TBD | Community-approved traditional knowledge |
| Code | 50M files | Canadian open-source projects |

| Capability | Use Case |
|---|---|
| Scientific Writing | Draft papers, translate to French, simplify for public |
| Grant Assistance | Format for Tri-Council, check eligibility, suggest reviewers |
| Data Discovery | Natural language queries to NORAI catalog |
| Code Generation | Scientific Python, R, Julia with Canadian library awareness |
| Regulation Lookup | PIPEDA, TCPS2, provincial health acts |
| Bilingual Reasoning | Seamless EN/FR switching mid-conversation |

| Phase | Cluster | GPUs | Time | Cost |
|---|---|---|---|---|
| Pre-training | Narval + Béluga | 512 A100s | 3 months | $8M |
| Fine-tuning | Cedar | 128 A100s | 2 weeks | $500K |
| RLHF | OVHcloud | 64 H100s | 1 month | $1M |

| Concern | Mitigation |
|---|---|
| Hallucination | RAG-first architecture, citation requirements |
| Bias | Canadian EDI training data curation, red-teaming |
| Misuse | Rate limiting, use case restrictions, audit logging |
| IP | Training data licensing tracked, opt-out respected |

| Aspect | Specification |
|---|---|
| Data Scale | 10M+ patient records across all provinces |
| Privacy Model | Federated learning + differential privacy ($\varepsilon$=1.0) |
| Network Topology | Hub-and-spoke with provincial secure enclaves |
| Compliance | PIPEDA, provincial health acts, TCPS2, OCAP |

**Network Architecture:**

**Genomic Analysis Pipelines:**

| Pipeline | Purpose | Privacy Approach |
|---|---|---|
| GWAS | Genome-wide association studies | Secure aggregation |
| Variant Calling | Identify cancer mutations | Local processing, summary stats only |
| Polygenic Risk | Multi-gene risk scoring | Federated model training |
| Pharmacogenomics | Drug response prediction | Differential privacy |

**Researcher Interface:** Researchers define pan-Canadian studies with specified provinces, privacy budgets, and ethics approvals. Analyses run locally at each province, and results return as aggregated statistics without individual patient data leaving provincial jurisdiction.

### 5.6 Borealis-Scholar AI Research Agent

**Objective:** Open-source, bilingual AI agent for superhuman literature search and synthesis

**Agent Architecture:**

**Capabilities:**

**API:** The Borealis-Scholar API provides deep literature search across multiple databases (Érudit, Semantic Scholar, PubMed, NORAI Catalog), synthesis of findings into literature review sections in either language, and identification of research gaps based on user queries.

### 5.7 Quantum-Classical Science Bridge

**Objective:** First national platform integrating quantum processors with HPC for scientific workflows

**Supported Quantum Approaches:**

**Hybrid Workflow:** Researchers define materials discovery workflows specifying classical and quantum resources. The platform automatically routes quantum portions (e.g., VQE for molecular ground states) to appropriate quantum hardware while classical optimization runs on Alliance clusters.

**Platform Architecture:**

**Quantum Advantage Tracking:**

| **Aspect** | **Specification** |
|---|---|
| Architecture | Multi-agent system with specialized sub-agents |
| Corpus | 5M+ Canadian papers, 200M+ global (via APIs) |
| Languages | English, French (native), 50+ (via translation) |
| Accuracy | >95% precision on literature search (benchmarked vs. PhD researchers) |

| **Agent** | **Role** | **Tools** |
|---|---|---|
| Orchestrator | Query understanding, task planning | LLM reasoning |
| Search Agent | Multi-database literature search | Semantic Scholar, PubMed, Érudit, NORAI Catalog |
| Synthesis Agent | Summarize, compare, critique papers | LLM + citation graph |
| Citation Agent | Verify claims, check retractions | CrossRef, Retraction Watch |
| Writing Agent | Draft literature reviews, abstracts | LLM + style templates |

| **Task** | **Time (Manual)** | **Time (Borealis-Scholar)** |
|---|---|---|
| Systematic literature search | 2-4 weeks | 2-4 hours |
| Literature review draft | 1-2 weeks | 1-2 hours |
| Citation verification | Days | Minutes |
| Gap analysis | Weeks | Hours |
| Translation (EN/FR) | Days | Seconds |

| **Aspect** | **Specification** |
|---|---|
| Quantum Hardware | Xanadu (photonic), D-Wave (annealing), IBM Q (via partnership) |
| Classical HPC | Alliance clusters (Narval, Cedar, Niagara) |
| Hybrid Workflows | Variational algorithms, quantum simulation, optimization |
| Access Model | Unified API, seamless job routing |

| **Approach** | **Hardware** | **Use Cases** |
|---|---|---|
| Photonic quantum | Xanadu Borealis | Molecular simulation, Gaussian boson sampling |
| Quantum annealing | D-Wave Advantage | Optimization, scheduling, materials |
| Gate-based | IBM Q (partnership) | Algorithm development, error correction research |
| Simulation | Alliance GPUs | Large-scale quantum simulation fallback |

| **Application** | **Classical Baseline** | **Quantum Target** | **Status** |
|---|---|---|---|
| Molecular ground states | Hours (CCSD(T)) | Minutes (VQE) | Approaching |
| Combinatorial optimization | NP-hard limits | Speedup on specific instances | Demonstrated |
| Machine learning kernels | $O(n^2)$ | O(log n) potential | Research |
| Materials simulation | Limited accuracy | Chemical accuracy | Target |

# 6. Technical Architecture

## 6.1 System Architecture Overview

## 6.2 Data Flow Architecture

## 6.3 Security Architecture

## 6.4 API Architecture

| Endpoint | Purpose | Authentication |
|---|---|---|
| `/v1/discover` | Search federated catalog | OAuth 2.0 + institutional IdP |
| `/v1/compute` | Submit federated computation | mTLS + consent tokens |
| `/v1/models` | Borealis model inference | API keys + rate limiting |
| `/v1/consent` | Manage data access consents | Role-based + community approval |
| `/v1/workflows` | Scientific workflow execution | Institutional + audit logging |

## 6.5 Technology Stack

| Layer | Technologies | Canadian Providers |
|---|---|---|
| **Compute** | Kubernetes, Slurm, Ray | Digital Research Alliance |
| **Storage** | Delta Lake, MinIO, Ceph | OVHcloud, IBM Canada |
| **ML/AI** | PyTorch, JAX, vLLM | Cohere, Untether AI |
| **Security** | Vault, SPIFFE/SPIRE | BlackBerry, Arctic Wolf |
| **Networking** | SCION, Tailscale | CANARIE, Cybera |
| **Observability** | Prometheus, Grafana | Canadian SaaS providers |

# 7. Five-Year Roadmap (2026-2030)

## 7.1 Phase 1: Foundation (2026)

**Objective:** Prove the model with working prototypes

| Quarter | Milestone | Deliverables |
|---|---|---|
| Q1 2026 | Platform foundation | Core infrastructure, team assembly |
| Q2 2026 | Computome MVP | Alliance cluster integration, dashboard |
| Q3 2026 | Federated Catalog MVP | Search across 3 federal departments |
| Q4 2026 | OCAP Pilot | Consent-as-code with 2 Indigenous communities |

**Year 1 Metrics:**

- 180 direct jobs created
- 3 federal data partnerships
- 2 Indigenous community pilots
- Computome serving 100+ researchers

**Budget:** $280M (see Section 11 for breakdown)

## 7.2 Phase 2: Expansion (2027)

**Objective:** Scale to national reach

| Quarter | Milestone | Deliverables |
|---|---|---|
| Q1 2027 | Provincial integration | BC, Alberta, Ontario health data protocols |
| Q2 2027 | Borealis-8B release | First Canadian science foundation model |
| Q3 2027 | Critical Minerals pilot | AI-generated prospect maps for NRCan |
| Q4 2027 | PermafrostGPT v1 | 1-km resolution forecasts for NWT |

**Year 2 Metrics:**

- 300 direct jobs
- 10 federal/provincial data partnerships
- 10 Indigenous community integrations
- 1,000+ active researchers

**Budget:** $280M

## 7.3 Phase 3: Integration (2028)

**Objective:** Full cross-provincial federation

| Quarter | Milestone | Deliverables |
|---|---|---|
| Q1 2028 | Pan-Canadian Health | Federated cancer research across 6 provinces |
| Q2 2028 | Borealis-70B | Multi-modal scientific reasoning |
| Q3 2028 | Climate Twin launch | National Arctic/boreal forecasting |
| Q4 2028 | Lab automation | TRIUMF, CLS integration |

**Year 3 Metrics:**

- 480 direct jobs

- All provinces participating
- 50+ Indigenous communities
- 5,000+ active researchers

**Budget:** $280M

### 7.4 Phase 4: Acceleration (2029)

**Objective:** Autonomous science capabilities

| Quarter | Milestone | Deliverables |
|---|---|---|
| Q1 2029 | MapleFusion | Fusion materials discovery pipeline |
| Q2 2029 | Borealis-405B | Hypothesis generation, experiment design |
| Q3 2029 | International pilots | Five Eyes, EU partnerships |
| Q4 2029 | Commercial APIs | Revenue-generating services |

**Year 4 Metrics:**

- 700 direct jobs
- First international partners
- $50M revenue
- 10,000+ researchers

**Budget:** $280M

### 7.5 Phase 5: Leadership (2030)

**Objective:** Global AI-for-science leadership

| Quarter | Milestone | Deliverables |
|---|---|---|
| Q1 2030 | Full autonomy | Closed-loop experiments |
| Q2 2030 | Export platform | Licensing to allied nations |
| Q3 2030 | Next-gen models | Borealis-1T planning |
| Q4 2030 | Sustainability | Self-funding operational model |

**Year 5 Metrics:**

- 1,000+ direct jobs
- $150M+ revenue
- 20,000+ researchers
- Global export presence

**Budget:** $280M

### 7.6 Total 5-Year Investment: $1.4B

### 7.7 ROI Projections

**Return on investment: 23x**

| Value Category | 5-Year Impact |
|---|---|
| Research productivity | $2B (time savings) |
| Critical minerals discovery | $15B (export potential) |
| Health outcomes | $5B (healthcare savings) |
| Climate adaptation | $10B (disaster cost reduction) |
| Talent retention | $1B (avoided brain drain) |
| **Total projected value** | **$33B+** |

# 8. Data Governance and Sovereignty

## 8.1 Principles

1. **Data Residency:** All data remains within Canadian jurisdiction
2. **Indigenous Data Sovereignty:** OCAP principles (Ownership, Control, Access, Possession)
3. **Consent-as-Code:** Machine-readable consent tokens govern all data access
4. **Minimal Disclosure:** Differential privacy and secure computation by default
5. **Audit Trail:** Immutable logs of all data access and computation

## 8.2 OCAP Implementation

| Principle | NORAI Implementation |
|---|---|
| **Ownership** | Community data portals, community-controlled access |
| **Control** | Consent-as-code policies, revocable permissions |
| **Access** | Community approval workflows, benefit sharing |
| **Possession** | Data never leaves community infrastructure without consent |

## 8.3 Compliance Framework

| Regulation | Approach |
|---|---|
| PIPEDA | Privacy-by-design, consent management |
| Provincial Health Acts | Federated computation, no data movement |
| OCAP | Indigenous community governance integration |
| ISED Security | Canadian-controlled infrastructure only |
| CCCS Guidelines | Threat intelligence, incident response |

# 9. Open Science Integration

## 9.1 The Problem: Fragmented Data Landscape

Canada has invested billions in open science platforms, but they remain siloed. A researcher studying climate impacts on health must navigate:

- Different authentication systems for each platform
- Incompatible data formats and metadata schemas
- No unified way to feed data into AI/ML pipelines
- Manual data wrangling consuming 40-60\

**NORAI solves this with a unified API wrapper** that provides a single interface to all Canadian open science data, optimized for AI model training and inference.

## 9.2 Canadian Open Science Platforms Inventory

Canada has over 30 major open science platforms across domains. NORAI will integrate them all:

| Platform | Data Holdings | Existing API | NORAI Wrapper |
|---|---|---|---|
| **CONP Portal** | 60+ datasets, 70+ tools, 17 institutions | DataLad, REST | `norai.neuro.conp` |
| **C-BIG Repository (MNI)** | Neurological patient data, imaging, biospecimens | LORIS API | `norai.neuro.cbig` |
| **CBRAIN** | Distributed neuroimaging compute | REST API | `norai.compute.cbrain` |
| **Brain-CODE** | Ontario Brain Institute data | Custom | `norai.neuro.braincode` |
| **NeuroLibre** | Reproducible publications | REST | `norai.neuro.neurolibre` |

**Neuroscience & Brain Research**

| Platform | Data Holdings | Existing API | NORAI Wrapper |
|---|---|---|---|
| **CanDIG** | Federated genomic analysis | GA4GH APIs | `norai.genomics.candig` |
| **CP4GPH** | Genomics & precision health ($14.5M platform) | GA4GH | `norai.genomics.cp4gph` |
| **Genome Canada** | 100,000+ Canadian genomes | Custom | `norai.genomics.gc` |
| **IHEC Data Portal** | 7,000+ epigenomic datasets | REST | `norai.genomics.ihec` |

**Genomics & Precision Health**

| Platform | Data Holdings | Existing API | NORAI Wrapper |
|---|---|---|---|
| **HDRN Canada** | National dataset registry | REST | `norai.health.hdrn` |
| **PopulationDataBC** | 5M residents, 40+ years longitudinal | Secure API | `norai.health.popdata` |
| **T-CAIREM** | AI-ready Ontario health data | REST | `norai.health.tcairem` |
| **Health Data Nexus** | | | |
| **ICES** | Ontario health services data | Secure API | `norai.health.ices` |

**Health Data Networks**

| Platform | Data Holdings | Existing API | NORAI Wrapper |
|---|---|---|---|
| **Ocean Networks Canada (Oceans 3.0)** | 14,000+ ocean sensors, petabytes of data | REST, Python SDK | `norai.ocean.onc` |
| **OBIS Canada** | Marine biodiversity, species distribution | REST | `norai.ocean.obis` |
| **DFO Open Data** | Fisheries, oceanographic conditions | CKAN | `norai.ocean.dfo` |
| **SLGO (St. Lawrence Global Observatory)** | St. Lawrence ecosystem data | REST | `norai.ocean.slgo` |

**Ocean & Marine Science**

| Platform | Data Holdings | Existing API | NORAI Wrapper |
|---|---|---|---|
| **OSDP (NRCan/ECCC)** | Environmental assessment, cumulative effects | REST | `norai.env.osdp` |
| **CANUE** | Urban environmental exposures, pan-Canadian | REST | `norai.env.canue` |
| **ECCC Climate Data** | Historical weather, climate projections | REST | `norai.env.climate` |
| **FOSRC** | Federal scientific publications | OAI-PMH | `norai.env.fosrc` |

**Climate & Environment**

| Platform | Data Holdings | Existing API | NORAI Wrapper |
|---|---|---|---|
| **Polar Data Catalogue** | 3,000+ datasets, 3M files, RADARSAT imagery | REST | `norai.arctic.pdc` |
| **CCIN** | Cryospheric data, permafrost, sea ice | REST | `norai.arctic.ccin` |
| **ArcticNet** | Integrated Arctic monitoring | Custom | `norai.arctic.arcticnet` |

**Arctic & Polar Research**

**Physics & Big Science Facilities**

**Biodiversity & Ecology**

**Agriculture & Food**

**General Research Infrastructure**

## 9.3 The NORAI Data Gateway: Unified API Wrapper

NORAI provides a single API wrapper that abstracts the complexity of 30+ platforms into a consistent interface optimized for AI/ML workflows:

## 9.4 API Wrapper Specifications

**Core Endpoints**

**Data Format Conversions**

| Platform | Data Holdings | Existing API | NORAI Wrapper |
|---|---|---|---|
| **Canadian Light Source (XASDB)** | X-ray spectroscopy, 1000+ spectra | REST | `norai.physics.cls` |
| **TRIUMF Tier-1** | Particle physics, ATLAS experiment | Grid | `norai.physics.triumf` |
| **CADC** | 2+ petabytes astronomy data, Hubble, JWST | TAP, REST | `norai.physics.cadc` |
| **SNOLAB** | Dark matter, neutrino experiments | Custom | `norai.physics.snolab` |

| Platform | Data Holdings | Existing API | NORAI Wrapper |
|---|---|---|---|
| **NatureCounts** | 200M+ bird observations | REST | `norai.bio.naturecounts` |
| **NatureServe Canada** | Species at risk, conservation status | REST | `norai.bio.natureserve` |
| **GBIF Canada** | Global biodiversity data | REST | `norai.bio.gbif` |
| **WildTrax (ABMI)** | Wildlife monitoring, acoustic data | REST | `norai.bio.wildtrax` |

| Platform | Data Holdings | Existing API | NORAI Wrapper |
|---|---|---|---|
| **AAFC Open Data** | Agricultural science, crop data | CKAN | `norai.agri.aafc` |
| **Canadian Soil Information Service** | National soil database | REST | `norai.agri.cansis` |

| Platform | Data Holdings | Existing API | NORAI Wrapper |
|---|---|---|---|
| **FRDR** | All Canadian research data | REST | `norai.data.frdr` |
| **Odesi** | 5,700+ social science datasets | REST | `norai.data.odesi` |
| **Dataverse Canada** | University research repositories | REST | `norai.data.dataverse` |
| **Scholars Portal** | Academic resources | REST | `norai.data.scholars` |

| Endpoint | Purpose | Example |
|---|---|---|
| gw.discover() | Federated search across all platforms | gw.discover("brain tumor MRI") |
| gw.get() | Retrieve dataset by ID | gw.get("conp:preventad-open") |
| gw.stream() | Stream large datasets in chunks | gw.stream(ds, chunk_mb=100) |
| gw.to_dataloader() | PyTorch/TensorFlow integration | gw.to_dataloader(ds, batch=32) |
| gw.export() | Export to ML-ready formats | gw.export(ds, format="parquet") |
| gw.cite() | Auto-generate citations | gw.cite(ds, style="apa") |

| Source Format | Target Formats | Use Case |
|---|---|---|
| DICOM, NIfTI | HDF5, Zarr, NumPy | Medical imaging to ML training |
| NetCDF, GRIB | Zarr, Parquet | Climate data to model input |
| VCF, BAM | Parquet, Arrow | Genomic data to ML pipelines |
| CSV, JSON | Parquet, Arrow | Tabular data to efficient loading |
| FITS | HDF5, Zarr | Astronomy to ML training |

**Authentication & Consent** The NORAI Gateway supports federated authentication via institutional identity providers (SAML/OIDC) and consent-aware data access. When datasets require consent (e.g., Indigenous health data), the system automatically initiates the appropriate consent workflow before granting access.

### 9.5 Platform Adapter Examples

The NORAI Gateway simplifies access to Canadian open science platforms:

- **Ocean Networks Canada:** Instead of platform-specific API calls with custom filters, researchers use a unified query interface that returns data in their preferred format (DataFrame, DataLoader for ML)
- **CONP Portal:** Instead of requiring DataLad and git-annex setup, researchers discover neuroscience datasets with simple queries and stream directly into training pipelines
- **All Platforms:** Consistent authentication, format conversion, and ML integration across 30+ data sources

### 9.6 Data Propagation & Synchronization

NORAI enables seamless data propagation across the country:

### 9.7 Benefits for AI Model Development

| Challenge | Current State | With NORAI Gateway |
|---|---|---|
| **Data discovery** | Search 30+ platforms separately | Single query across all platforms |
| **Authentication** | 30+ different accounts/tokens | One NORAI token, federated auth |
| **Data formats** | Manual conversion scripts | Automatic format conversion |
| **ML integration** | Custom data loaders per source | Native PyTorch/TensorFlow integration |
| **Consent tracking** | Manual paperwork | Automated consent-as-code |
| **Citation** | Manual bibliography management | Auto-generated citations |
| **Data updates** | Manual re-download | Event-driven propagation |
| **Reproducibility** | Ad-hoc versioning | Full provenance tracking |

### 9.8 Implementation Roadmap

| Phase | Timeline | Platforms Integrated | Capabilities |
|---|---|---|---|
| **Phase 1** | Q2 2026 | CONP, FRDR, ONC, PDC | Discovery, basic access |
| **Phase 2** | Q4 2026 | CanDIG, CADC, OSDP, DFO | Streaming, format conversion |
| **Phase 3** | Q2 2027 | Health platforms (with privacy) | Federated learning support |
| **Phase 4** | Q4 2027 | All 30+ platforms | Full propagation, ML pipelines |

### 9.9 Open Source Commitment

The NORAI Data Gateway will be released as open source (Apache 2.0), enabling:

- Platform operators to contribute adapters
- Researchers to self-host for sensitive data
- International partners to deploy compatible gateways

- Community-driven improvements

# 10. Canadian AI Ecosystem

## 10.1 National AI Institutes (Proposed Partnerships)

| Institution | Contribution | Location |
|---|---|---|
| Mila | Deep learning research, talent pipeline | Montreal |
| Vector Institute | Healthcare AI expertise | Toronto |
| Amii | Reinforcement learning, industry partnerships | Edmonton |

## 10.2 National Compute Infrastructure (Proposed Partnerships)

| Organization | Contribution | Coverage |
|---|---|---|
| Digital Research Alliance | HPC coordination | National |
| Compute Ontario | Regional HPC | Ontario |
| Calcul Quebec | Advanced computing | Quebec |
| ACENET | HPC services | Atlantic |
| WestGrid/BC DRI | HPC resources | Western |

## 10.3 Research Universities (Proposed Partnerships)

| University | Research Focus |
|---|---|
| University of Waterloo | Quantum computing, cybersecurity |
| University of Toronto | Medical AI, robotics |
| McGill University | Neuroscience-inspired AI, climate |
| University of Alberta | Reinforcement learning, NLP |
| University of British Columbia | Computer vision, sustainable AI |
| Université de Montréal | Deep learning foundations |
| University of Calgary | Energy, geoscience AI |
| Dalhousie University | Ocean science, health |

## 10.4 Federal Research Organizations (Proposed Partnerships)

## 10.5 Industry Ecosystem

| Organization | Contribution |
|---|---|
| National Research Council | Applied research, technology transfer |
| Natural Resources Canada | Geological data, climate datasets |
| Environment and Climate Change Canada | Weather, environmental monitoring |
| Canadian Space Agency | Earth observation, satellite data |
| National Defence / DRDC | Security applications |

| Sector | Canadian Companies |
|---|---|
| AI/ML | Cohere, Sanctuary AI, Untether AI, Element AI alumni |
| Cloud | OVHcloud Montreal, IBM Canada, TELUS |
| Cybersecurity | BlackBerry/Cylance, Arctic Wolf, eSentire |
| Quantum | Xanadu, D-Wave, Photonic Inc |
| Health Tech | WELL Health, Think Research, MedStack |
| Geospatial | MDA, Ecopia AI, Geotab |
| Indigenous Tech | Animikii, First Nations Technology Council |

# 11. Multi-Stakeholder Alignment

## 11.1 Federal Stakeholder Mapping

| Stakeholder | Primary Interest | NORAI Alignment |
|---|---|---|
| **ISED** | Innovation, AI strategy, sovereignty | Canadian AI leadership, job creation |
| **NRCan** | Resources, energy, northern development | Critical minerals, permafrost, energy AI |
| **ECCC** | Climate, environment | Borealis Climate Twin, emissions modeling |
| **PMO/PCO** | Strategic positioning, reconciliation | Sovereignty narrative, Indigenous partnership |
| **Health Canada** | Health innovation | Privacy-preserving health AI |
| **Transport Canada** | Infrastructure resilience | Permafrost impacts, northern logistics |
| **DND/CAF** | Arctic security, strategic resources | Northern infrastructure, resource security |
| **CSE/CCCS** | Cybersecurity, data protection | Sovereign infrastructure, threat intelligence |
| **NRC** | Research translation | Commercialization pathway |
| **CIRNAC** | Indigenous relations | OCAP implementation, community benefits |
| **StatsCan** | National data | Federated access, privacy preservation |
| **DRDC** | Defence R&D | Dual-use applications |

## 11.2 Provincial Alignment

## 11.3 Indigenous Partnership Framework

| Province | Key Interest | NORAI Offering |
|---|---|---|
| Quebec | AI leadership, French language | Bilingual platform, Mila partnership |
| Ontario | Health data, tech sector | Vector partnership, health federation |
| Alberta | Energy, resources | Energy AI, Amii partnership |
| BC | Climate, health | Climate twin, health federation |
| Northern territories | Infrastructure, sovereignty | PermafrostGPT, Arctic focus |
| Atlantic | Ocean science, health | Ocean data, regional HPC |

| Partner Type | Engagement Model |
|---|---|
| National organizations (ITK, AFN, MNC) | Governance advisory, policy co-development |
| Regional organizations | Data sharing protocols, benefit agreements |
| Individual communities | Consent-as-code pilots, capacity building |
| Indigenous tech companies | Contracting, technology partnerships |

# 12. Funding and Economic Impact

## 12.1 Federal Budget Context (2024-2025)

The Canadian government has committed significant funding to AI and science infrastructure. NORAI aligns with and extends these existing commitments:

| Federal Program | Commitment (5 years) | Source |
|---|---|---|
| AI Sovereign Compute Strategy | $2,000M | Budget 2024 |
| Tri-Council Grants (CIHR/NSERC/SSHRC) | $1,800M | Budget 2024 |
| Public AI Infrastructure | $926M | Budget 2025 |
| Student Scholarships | $825M | Budget 2024 |
| Quantum Computing Infrastructure | $334M | Budget 2025 |
| AI Compute Access Fund | $300M | Budget 2024 |
| Digital Research Alliance | $285M | FES 2024 |
| **Total Available** | **$6,470M** | |

## 12.2 NORAI Funding Alignment

NORAI requests $1.4B over 5 years - approximately **21% of new AI/science commitments** - allocated from existing and planned federal programs:

## 12.3 Annual Budget Allocation

## 12.4 Cost Efficiency Analysis

## 12.5 Economic Multiplier Effects

*Note: The $33B projected value includes downstream economic impacts from research outcomes (critical minerals, health savings, climate resilience, etc.)*

## 12.6 Job Creation by Region

## 12.7 Revenue Model (Post-2028)

*Revenue contributes to long-term sustainability, reducing ongoing federal funding requirements by 50% by 2032.*

| Source | NORAI Allocation | Rationale |
|---|---|---|
| Public AI Infrastructure ($926M) | $500M | NORAI is exactly the sovereign AI infrastructure this fund targets |
| New Tri-Council Allocation ($1.8B) | $350M | 19% of new grants to accelerate research productivity by 40% |
| Digital Research Alliance ($285M) | $200M | Mandate extension to include AI orchestration and governance |
| AI Sovereign Strategy ($2B) | $200M | 10% for platform connecting compute infrastructure to researchers |
| Provincial Contributions | $100M | QC, ON, AB, BC contribute based on researcher utilization |
| Industry Partnerships | $50M | Cohere, D-Wave, Xanadu, BlackBerry in-kind and cash |
| **Total NORAI** | **$1,400M** | **21% of commitments → 23x ROI** |

| Category | 2026 | 2027 | 2028 | 2029 | 2030 | Total |
|---|---|---|---|---|---|---|
| Infrastructure | $80M | $70M | $60M | $50M | $40M | $300M |
| Software Development | $60M | $70M | $80M | $70M | $60M | $340M |
| Research & Prototypes | $70M | $60M | $50M | $60M | $70M | $310M |
| Personnel & Operations | $50M | $60M | $70M | $80M | $90M | $350M |
| Partnerships | $20M | $20M | $20M | $20M | $20M | $100M |
| **Annual Total** | **$280M** | **$280M** | **$280M** | **$280M** | **$280M** | **$1,400M** |

| Metric | Traditional Grants | Compute Canada | NORAI (Year 1) | NORAI (Year 5) |
|---|---|---|---|---|
| Cost per active researcher | $85,000 | $45,000 | $56,000 | $18,700 |
| Research productivity gain | Baseline | +15% | +25% | +50% |
| ROI multiplier | 4x | 6x | 8x | 23x |

| Investment | Direct Value | Multiplier Effect | Total Value |
|---|---|---|---|
| Compute infrastructure | $300M | 3x | $900M |
| Software development | $340M | 4x | $1.36B |
| Research programs | $310M | 5x | $1.55B |
| Personnel & operations | $350M | 3x | $1.05B |
| Partnerships | $100M | 6x | $600M |
| **Total** | **$1.4B** | **4x avg** | **$5.5B direct** |

| Region | Year 1 | Year 3 | Year 5 |
|---|---|---|---|
| Quebec | 60 | 150 | 300 |
| Ontario | 50 | 140 | 280 |
| British Columbia | 25 | 60 | 150 |
| Alberta | 25 | 60 | 120 |
| Atlantic | 15 | 50 | 100 |
| Northern/Territories | 5 | 20 | 50 |
| **Total** | **180** | **480** | **1,000** |

| Revenue Stream | Year 3 | Year 5 |
|---|---|---|
| Commercial API licenses | $10M | $50M |
| Industry partnership fees | $15M | $40M |
| International licensing (Five Eyes, EU) | $5M | $30M |
| Training & certification | $5M | $20M |
| Consulting services | $5M | $10M |
| **Total** | **$40M** | **$150M** |

# 13. Why NORAI

## 13.1 Differentiators vs. US Genesis

| Capability | Genesis (US) | NORAI (Canada) |
|---|---|---|
| Indigenous governance | None | OCAP-native (world first) |
| Bilingual platform | No | EN/FR/Inuktitut |
| Provincial federation | N/A | 13 jurisdictions |
| Export to allies | Unlikely | Five Eyes, EU ready |
| Open science | DOE-controlled | FAIR principles |
| Data sovereignty | US jurisdiction | 100% Canadian |

## 13.2 Canadian Competitive Advantages

1. **OCAP/Indigenous governance** - Unmatched globally, export potential
2. **Bilingual by design** - Exportable to Francophone nations, UN agencies
3. **Privacy-preserving federation** - Provincial health data accessible without centralization
4. **Allied trust** - Five Eyes, EU prefer Canadian over US platforms for sensitive data
5. **Talent pool** - World-class AI researchers (Bengio, Hinton, Sutton alumni)

## 13.3 Risk Mitigation Summary

| Risk | Mitigation |
|---|---|
| US cloud dependency | Canadian-first infrastructure strategy |
| Talent competition | Competitive salaries, compelling mission |
| Provincial coordination | Federal coordination office, incentive alignment |
| Indigenous trust | Co-development from day one, benefit sharing |
| Technology obsolescence | Modular architecture, continuous modernization |

# 14. Risk Assessment and Mitigation

## 14.1 Competitive Landscape Analysis

**Cloud Platform Alternatives** NORAI must differentiate from existing options. Here's how we compare:

| **Capability** | **AWS** | **Azure** | **Google** | **NORAI** |
|---|---|---|---|---|
| Canadian sovereignty | US juris. | US juris. | US juris. | **100% CDN** |
| Data residency | CDN region | CDN region | Limited | **Guaranteed** |
| Indigenous governance | None | None | None | **OCAP** |
| Science-specific tools | Generic | Generic | Some ML | **Purpose-built** |
| Bilingual support | Limited | Limited | Limited | **EN/FR** |
| Research integration | None | None | None | **Built-in** |
| Cost for researchers | Commercial | Commercial | Commercial | **Subsidized** |

**Why researchers won't just use AWS/Azure:**

- US cloud providers subject to CLOUD Act (US government data access)
- No integration with Canadian research workflows (CCDB, Alliance accounts)
- No OCAP compliance mechanisms
- Generic tools require significant customization for science
- Commercial pricing prohibitive for most research budgets

| **Initiative** | **Threat Level** | **NORAI Response** |
|---|---|---|
| US Genesis | High | Build what Genesis cannot: OCAP, bilingual, federation |
| EU AI4Science | Medium | Collaborate via Five Eyes/EU channels, not compete |
| UK AI Research | Low | Potential partnership, similar values |
| China AI platforms | N/A | Non-aligned; NORAI serves as allied alternative |

**International AI-for-Science Competition**

## 14.2 Technology Risks

| **Risk** | **Probability** | **Impact** | **Mitigation** |
|---|---|---|---|
| Borealis models underperform vs commercial LLMs | Medium | High | Hybrid approach: Borealis for science-specific tasks, integration APIs for general LLMs |
| Quantum computing disrupts current approaches | Low (5yr) | High | Partnership with Xanadu/D-Wave; quantum-ready architecture |
| Foundation model costs continue to decline | High | Medium | Focus on federation and governance (defensible moats), not raw compute |
| Talent acquisition fails | Medium | High | Competitive comp, remote-first, mission-driven recruitment |
| Security breach | Low | Very High | CSE partnership, regular audits, bug bounty program |

## 14.3 Execution Risks

| Risk | Probability | Impact | Mitigation |
|---|---|---|---|
| Delayed federal funding approval | High | High | Phase deliverables to show value with $50M initial tranche |
| Provincial governments don't participate | Medium | High | Start with willing provinces (QC, ON, BC), demonstrate value, expand |
| Universities reluctant to share data | High | Medium | Federated approach (data stays in place), clear value proposition |
| Indigenous communities distrust initiative | Medium | Very High | Co-design from day one, FNIGC partnership, benefit sharing |
| Key personnel departure | Medium | High | Distributed leadership, documentation, competitive retention |
| Scope creep | High | Medium | Strong program governance, MVP-first approach |

## 14.4 Governance Structure

To ensure accountability and multi-stakeholder alignment, NORAI proposes the following governance:

## 14.5 Intellectual Property Strategy

| IP Type | Ownership | Rationale |
|---|---|---|
| Core platform code | Open source (Apache 2.0) | Maximize adoption, enable export |
| Borealis model weights | Canadian government | Sovereign asset, controlled licensing |
| Research outputs | Researchers/institutions | Standard academic norms |
| Indigenous data derivatives | Communities | OCAP compliance |
| Commercial applications | Licensees | Revenue generation |

**Export Control:**

- Borealis models classified as controlled technology
- Licensing to Five Eyes partners, EU under bilateral agreements
- No export to non-allied nations without Cabinet approval

## 14.6 Security & Threat Model

| Threat | Likelihood | Mitigation |
|---|---|---|
| Nation-state cyber attack | High | CCCS partnership, air-gapped sensitive systems |
| Insider threat | Medium | Clearances for key personnel, audit logging |
| Supply chain compromise | Medium | Canadian-sourced hardware where possible, code audits |
| Data exfiltration | Low | Encryption at rest/transit, DLP tools |
| Service disruption | Medium | Multi-region redundancy, incident response plan |

**Security Certifications Target:**

- CCCS IT Security Assessment (Year 1)

- SOC 2 Type II (Year 2)
- ISO 27001 (Year 2)
- Protected B capability (Year 3)

# 15. Addressing Stakeholder Concerns

## 15.1 Provincial Governments

**Concern:** "Why should we contribute when the federal government controls everything?"

| Issue | NORAI Response |
|---|---|
| Control | Provincial representatives on Steering Committee and Advisory Council |
| Data sovereignty | Data physically remains in province; only metadata federated |
| Funding fairness | Contributions proportional to researcher utilization; benefits flow back |
| Health data | No centralization; federated learning preserves provincial jurisdiction |
| Jobs | Regional hiring requirements; proportional job distribution |

**Provincial Value Proposition:**

| Province | Specific Benefit |
|---|---|
| Quebec | Bilingual platform validates Mila investment; FR-first science AI |
| Ontario | Vector partnership amplified; health AI leadership |
| Alberta | Energy transition AI; Amii collaboration |
| BC | Climate/ocean science; Vancouver tech ecosystem integration |
| Atlantic | Ocean science hub; distributed jobs beyond Toronto/Montreal |
| Northern | PermafrostGPT directly serves territorial needs; Arctic sovereignty |

## 15.2 Research Universities

**Concern:** "We've seen government IT projects fail. Why will this be different?"

| Issue | NORAI Response |
|---|---|
| Past failures | Agile delivery with quarterly milestones; public progress tracking |
| Control over research | NORAI is infrastructure, not research direction; researchers choose what to study |
| Data ownership | Researchers retain ownership; NORAI provides tools, not mandates |
| Overhead burden | Single sign-on with institutional credentials; self-service portal |
| Competitive advantage | Early adopters get priority access; collaboration tools |

**University Adoption Incentives:**

## 15.3 Industry Partners

**Concern:** "What's in it for us? This sounds like a government boondoggle."

**Industry Participation Models:**

**Industry ROI Projections:**

## 15.4 Indigenous Communities

**Concern:** "Governments have promised data sovereignty before and not delivered."

**Indigenous Engagement Phases:**

## 15.5 Federal Departments

**Concern:** "This duplicates existing investments."

**De-duplication Principle:** NORAI builds the connective tissue between existing investments, not parallel systems.

| Incentive | Details |
|---|---|
| Compute credits | $50K-$500K compute grants for early adopter institutions |
| Co-op positions | Funded student positions for platform development |
| Publication support | Citation tools, data provenance, reproducibility features |
| Grant competitiveness | NORAI usage integrated with Tri-Council applications |

| Issue | NORAI Response |
|---|---|
| Commercial viability | Revenue-generating APIs by Year 3; commercial licensing |
| Procurement fairness | Open RFPs; Canadian-first but competitive |
| IP protection | Clear IP frameworks; industry retains commercial IP |
| Government pace | Agile delivery; industry-standard development practices |
| Market access | International export platform; Five Eyes/EU market entry |

| Model | Description | Example Companies |
|---|---|---|
| Infrastructure partner | Provide/operate core infrastructure | OVHcloud, IBM Canada, TELUS |
| Technology partner | Contribute specific capabilities | Cohere (LLMs), Xanadu (quantum), BlackBerry (security) |
| Solution partner | Build applications on platform | Canadian SaaS companies, consultancies |
| Data partner | Contribute proprietary datasets | Mining companies, health systems |

| Partner Type | Year 1-2 | Year 3-5 |
|---|---|---|
| Infrastructure | $5-20M contracts | $10-40M contracts + commercial revenue share |
| Technology | $2-10M contracts | Platform revenue licensing (3-5% royalties) |
| Solution | Early access, co-development | Commercial market access |

| Issue | NORAI Response |
|---|---|
| Historical trust deficit | FNIGC and ITK partnership from governance design phase |
| OCAP implementation | Machine-enforceable consent; community veto power |
| Benefit sharing | Revenue sharing for commercial use of Indigenous data |
| Capacity building | Training programs, funded positions in Indigenous communities |
| Community control | Dedicated Indigenous Governance Circle with binding authority |

| Phase | Timeline | Activities |
|---|---|---|
| 1. Relationship building | Q1-Q2 2026 | Formal engagement with ITK, AFN, FNIGC; listening sessions |
| 2. Co-design | Q3-Q4 2026 | Joint design of consent-as-code; pilot community selection |
| 3. Pilot | 2027 | 2-3 community pilots with full OCAP compliance |
| 4. Expansion | 2028+ | Community-led expansion; Indigenous tech company partnerships |

| Department | Existing Investment | NORAI Augmentation |
|---|---|---|
| ISED | AI Sovereign Compute Strategy ($2B) | NORAI is the platform that makes compute useful for science |
| NRCan | Geological Survey digitization | Critical Minerals AI analyzes existing data |
| ECCC | Weather/climate data systems | Borealis Climate Twin adds prediction capability |
| Health Canada | Health data modernization | Federated learning enables analysis without centralization |
| DND | Arctic sovereignty investments | PermafrostGPT supports infrastructure planning |

# 16. Call to Action

## 16.1 The Choice

The Government of Canada has three realistic paths:

1. **Build from scratch** - 8-10 years, uncertain outcome, talent dispersal continues
2. **Depend on US platforms** - Sovereignty risk, data export, economic leakage
3. **Partner with Canadian industry** - Deliver results by 2027, jobs stay in Canada

**We choose door #3.**

## 16.2 Immediate Next Steps

1. **Establish NORAI Program Office** within ISED (Q1 2026)
2. **Secure Year 1 funding** ($280M from aligned federal programs - see Section 11)
3. **Formalize partnerships** with Digital Research Alliance, Mila, Vector, Amii
4. **Launch Indigenous consultation** with ITK, AFN, FNIGC
5. **Issue RFP** for Canadian technology partners

## 16.3 The Ask

We invite the Government of Canada to:

- Open discussions with Canadian AI companies and research institutions
- Establish a dedicated program office for sovereign AI infrastructure
- Commit Year 1 funding ($280M) aligned with existing federal AI/science commitments
- Set a target: **First federated Canadian science experiment by Canada Day 2026**

---

## 16.4 Contact

**Semanttica Web Technologies (SWT) Inc.** Montreal, Quebec, Canada

Website: norai.ca Email: hello@norai.ca

---

**Canada does not need to catch up. Canada needs to build what only Canada can build. With NORAI, Canada leads from the True North.**

---

## References


[1] Digital Research Alliance of Canada, "2024 Resource Allocations Competition Results," 2024. Available: https://www.alliancecan.ca/en/2024-resource-allocations-competition-results

[2] Government of Canada, "Budget 2024: Fairness for Every Generation," April 2024. Available: https://budget.canada.ca/2024/home-accueil-en.html

[3] Innovation, Science and Economic Development Canada, "Canadian Sovereign AI Compute Strategy," December 2024. Available: https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy

[4] U.S. Department of Energy, "Energy Department Launches Genesis Mission to Transform American Science and Innovation Through the AI Computing Revolution," November 2025. Available: https://www.energy.gov/articles/energy-department-launches-genesis-mission-transform-american-science-and-innovation

[5] NSERC, "Discovery Research Program Competition Results," 2024. Available: https://www.nserc-crsng.gc.ca/nserc-crsng/fundingdecisions-decisionsfinancement/researchgrants-subventionsderecherche/index_eng.asp

[6] Anaconda, "State of Data Science 2020," 2020. Available: https://www.anaconda.com/state-of-data-science-2020

[7] Canadian Institutes of Health Research, "Budget 2024: Message from the President," April 2024. Available: https://cihr-irsc.gc.ca/e/53894.html

[8] Government of Canada, "Government of Canada invests in 7,700 world-class researchers and projects across the country," March 2024. Available: https://www.canada.ca/en/innovation-science-economic-development/news/2024/03/government-of-canada-invests-in-7700-world-class-researchers-and-projects-across-the-co
html

[9] Argonne National Laboratory, "Energy department launches Genesis Mission to transform American science and innovation through the AI computing revolution," November 2025. Available: https://www.anl.gov/article/energy-department-launches-genesis-mission-transform-american-science-and-innovation-th

[10] NVIDIA, "NVIDIA and Oracle to Build US Department of Energy's Largest AI Supercomputer for Scientific Discovery," October 2025. Available: https://nvidianews.nvidia.com/news/nvidia-oracle-us-department-of-energy-ai-supercomputer-scientific-discovery

[11] Government of Canada, "What We Heard: Tri-agency engagement with the research community on modernization of the federal research support system," 2024. Available: https://www.canada.ca/en/research-coordinating-committee/services/publications/progress-reports/2024/what-we-heard-tri-agency-engagement-with-the-research-community-on-modernization-o
html

[12] Digital Research Alliance of Canada, "The Alliance applauds the Federal government's launch of the Canadian Sovereign AI Compute Strategy," December 2024. Available: https://www.alliancecan.ca/en/latest/news/alliance-applauds-federal-governments-launch-canadian-sovereign-ai-compute-strategy