

Responsible AI Infrastructure for Indigenous Research Communities

Emerson Sarmiento · Karla Pinzón · Lenna Stahl

May 2026 Network Presentation

OUR TEAM

Bard Graduate Programs in Sustainability · MBA in Sustainability



Emerson Sarmiento

PROJECT MANAGER

Bard MBA · May 2026

Concentration: Circular Value Chain Management



Karla Pinzón

MBA RESEARCHER

Bard MBA · Expected May 2027

Concentration: Circular Value Chain Management



Lenna Stahl

MBA RESEARCHER

Bard MBA · May 2026

Concentration: Impact Investing

Agenda

01 **The next layer of research**

02 **What communities actually need**

03 **What alternatives exist**

04 **Sustainability strategies**

05 **Trade Offs**

06 **A decision toolkit**

The Next Layer of Research

Phase 1 - Identifying alternatives:

Mapping the AI supply chain and the case for alternatives

Phase 2 - Evaluating them:

Evaluating which infrastructure alternatives are real and usable, and for whom

Infrastructure. Sustainability. Sovereignty. Action.

Questions we wanted to explore further

01 Do these alternatives actually exist in practice?

02 Who is building them and how mature are they?

03 Could communities with limited IT capacity realistically adopt them?

04 What tradeoffs does each option create?

To answer these questions, we reviewed 40+ sources and surveyed 5 pods

What Do Communities Actually Need?

Most pods don't need hyperscale compute.
That changes everything.

Five pods. Five infrastructure realities. One honest finding.

The toolkit has to serve five different situations, not one composite community. What the network actually needs is pluralism at the governance layer, not a single sovereign-server template.

Sovereignty has multiple layers

“The confidence barrier is less technical than institutional. The servers can be built. The legitimacy has to be negotiated.”

— *T'Karonto Pod*

Data layer

Where research files physically live. Sovereignty here means refusing to store cultural, ceremonial, or language data on infrastructure the community can't audit or revoke.

Compute layer

Where models actually run. Local hosting matters but ownership without governance, accountable staff, and institutional recognition does not constitute control.

Model layer (Hidden)

Whose pretrained models you depend on. Even fully offline pipelines typically begin by downloading a models trained on corpora the community had no role in shaping.

Governance layer

Who has legal standing to decide all of the above. Named entity, succession plan, and host-institution recognition for ethics, funding, and publication.

ML tasks $\neq$ frontier AI, and it matters a lot

Frontier AI assumption

- ✗ Massive GPU clusters
- ✗ Hyperscale cloud dependency
- ✗ Millions in infrastructure cost
- ✗ Enormous energy & water footprint
- ✗ Sovereignty nearly impossible

Community ML reality

- ✓ **Small AI on Raspberry Pi**
Haudenosaunee Pod runs and retrains small AI on edge devices
- ✓ **Language and cultural data work**
Anishinaabemowin game (Biskaabiiyaang), William Jones archive, Blackfoot ELAN work
- ✓ **Semantic mapping and topology**
Wihable Sa Center and T'Karonto sovereignty-compliant offline tools
- ✓ **Multi-layer fine-tuning stacks**
Hiringa Te Mahara designs, builds, and uses AI across 5 compute layers
- ✓ **Modest hardware footprint**
Sovereignty becomes achievable when scaled to the actual workload

What communities told us they prioritize

Sovereignty AND institutional standing

Ownership alone isn't enough. The hardware can be built; the legitimacy has to be negotiated. T'Karanto Pod names four conditions: reliable hosting, named governance entity with legal standing, community-accountable technical staff, and host-institution recognition for ethics, funding, and publication.

Affordability that outlasts launch funding

Every respondent named cost as a top barrier. Broadband case studies show what fails: one-time grants that don't cover operations. Niitsitapi Pod adds a governance dimension: elder-led decisions take time, and member availability is its own cost.

Low technical burden + transferable expertise

Lack of IT expertise is a universal barrier. Replication fails when knowledge is individual-dependent. Haudenosaunee Pod's response: a standardized hardware AND software package other communities can deploy locally, not just a framework.

Governance fit, not one architecture for all

A sovereign server hosted with one nation sets a precedent for others. What works for Winhablé S'a may not transfer cleanly to the Haudenosaunee or Niitsitapi pods. The toolkit must help pods reason about which model fits which relationship.

Cross-pod findings

Data sovereignty is near-universal, but control isn't

83% (5 of 6) pod respondents rated Indigenous data control a 5/5 priority. Only 1 respondent reported full control. The other 5 said "somewhat", meaning the aspiration is strong but the infrastructure doesn't match it yet.

Model supply chain is an unsolved sovereignty problem

Even a fully offline pipeline typically begins by downloading a pretrained model from the US or Europe, trained on corpora the community had no role in shaping. Infrastructure sovereignty addresses where work runs. Model sovereignty is a separate, unsolved problem.

Heterogeneity is the norm, not the exception

No two pods run the same stack. Hiringa Te Mahara uses 5 different compute layers. Haudenosaunee combines own server + Raspberry Pi + Arduino + SharePoint. Treating the network as one infrastructure problem flattens what's actually a set of different problems.

Environmental impact is a concern, not yet a plan

100% of respondents said environmental cost has come up but average awareness score was only 3.2/5. Niitsitapi pod specifically said, "We'd like to address it but don't know where to start."

Where each pod is today

Pod	Current compute	Data control	Sovereignty (1–5)	Top barrier
T'Karonto (Respondent 1)	Personal laptops + offline hard drives only	Somewhat	5	Cost + institutional lock-in
Hiringa Te Mahara	Laptops + university systems + own servers	Somewhat	5	Cost + technical complexity
T'Karonto (Respondent 2)	Laptops + university + some Indigenous-hosted	Somewhat	4	Cost + institutional requirements
Niitsitapi pod	Laptops + university + pod-owned devices	Somewhat	5	Cost + don't know alternatives
Haudenosaunee Pod	Own server + community land hosting + Raspberry Pi edge devices	Full control ✓	5	Cost + time
Wihanble Sa Center	Laptops + university + cloud (imposed)	Somewhat	5	Cost + technical complexity

★ Haudenosaunee Pod is the only pod in the network with community-controlled infrastructure already operational — and is actively building a replication model for others.

In their own words

“Our pod strikes a balance between community work and research/art as experimentation. Our goal is not to create tribal infrastructure but to test ideas in a controlled space.”

— Wihanble S’aat Pod, Bard College

“Placing data in community at the control of community limits the capacity of others. Think small-scale. The goal is to fight against large-scale access to Indigenous data.”

— Haudenosaunee Pod, Western University


“We are both urban and have a large non-urban in community project. The latter would love to see a sovereign stack on Anishnabe land. The former is interested on whether there could be a universities agreement for hosting on university land by treaty holders.”

— T’Karonto Pod, Respondent 2

“It would need to be Niitsitapi pod-led in partnership with the Blackfoot community and in accordance with OCAP principles in Blackfoot contexts.”

— Niitsitapi Pod, University of Lethbridge

Network coverage and data still to gather



Data still to gather

01 1 of 6 established pods unreachable

Ka Hōwā'i Pae 'Āina (UH at West O'ahu) has not responded to the questionnaire. The Pacific perspective is missing entirely from the current dataset.

02 No cost or environmental baselines

Cost was named as a barrier by every respondent and environmental impact by 100%, but we didn't collect actual numbers. Awareness average is 3.2/5.

03 Model supply chain unmapped

Respondent #1 named pretrained models as the deepest hidden dependency, but no systematic audit exists yet of which models the network actually depends on.

04 8 community partners unsurveyed

We reached pod leads and researchers. The 8 community organizations connected to the network have not yet been asked.

What Alternatives Actually Exist?

Not theory. Actual implementation.

We assessed alternatives on four dimensions:

- 01 Existence: Is it actually deployed somewhere, not just proposed?
- 02 Maturity: Ready now, emerging, or experimental?
- 03 Community fit: Can a pod with limited IT capacity operate it?
- 04 Replicability: What does it take for another pod to follow?

Infrastructure alternatives by maturity tier

Ready Now

- Refurbished server systems
- Air-cooled basic local hosting
- Renewable-powered co-location
- Small model deployment

Emerging

- Small modular data centers
- ODINN Omnia portable compute
- Waste heat recovery systems
- Edge computing networks
- Cooperative cloud infrastructure (ie. Nubo cooperative cloud)

Experimental

- Recycled/seawater cooling (SWAC)/ocean thermal air conditioning
- Distributed mesh compute (DePIN)
- Fully off-grid renewable compute

Indigenous communities are already building right now

Native BioData Consortium

Indigenous-controlled genomic research infrastructure with legal and technical controls over data access, use, and reuse.

→ *Sovereignty requires institutional + legal + technical controls together, not just ownership.*

Tribal Digital Village

Locally run connectivity built through grants, community labor, mountain-top towers, and solar-powered infrastructure.

→ *Ownership works when right-sized, locally governed, grant-supported, and tied to anchor institutions.*

Nubo Cooperative Cloud

Digital infrastructure governed through cooperative model → trust as an ownership and accountability feature, not a vendor promise.

→ *Cooperative or federation-based infrastructure is viable for storage, data stewardship, and lower-intensity workloads.*

Native Nations Broadband

Native nations building their own broadband when market providers fail → a practical analogy for technology self-determination.

→ *Ownership builds control and capacity, but requires staffing, maintenance, grants, and long-term sustainability.*

Sustainability Strategies

Structural Redesign vs. Superficial Mitigation

Superficial mitigation

Reduce the harm of existing infrastructure without changing it

Carbon credits, renewable PPAs → real strategies but the infrastructure stays the same

Structural redesign

Choose infrastructure whose footprint is smaller by design

Small models, cooperative compute, refurbished hardware, geothermal power

Strategy maturity matrix

Strategy	Maturity	Type	Research note
Carbon credits	Mature	Superficial mitigation	<i>Well-documented; limited systemic change</i>
Renewable PPAs	Mature	Partial structural	<i>Indigenous-owned renewables now 85% of new projects (Energy Mix, 2026)</i>
Onsite renewables	Moderate	Structural	<i>Geothermal emerging as firm clean power option (Stanford, 2025)</i>
Waste heat recovery	Emerging	Structural + Community	<i>Pilot cases in Appalachia; requires nearby heat users (ReImagine App., 2026)</i>
Air/liquid cooling	Mature	Structural	<i>Zero-water cooling viable; may raise energy use (Sustainability Mag., 2025)</i>
Refurbished servers	Mature (niche)	Structural + Circular	<i>Strong circularity case; vendors and networks exist</i>
Small models	Increasingly viable	Structural	<i>Best fit for language/cultural workloads; lower energy and cost</i>
Cooperative ownership	Emerging	Structural + Sovereignty	<i>Nubo & Platform.Coop models live; capital and uptime remain key risks</i>

Refurbished hardware: a values-aligned choice

**Extend.
Reuse.
Repair.**

Alipio et al. (2024): sustainable data centers should be designed around existing watershed and climate systems – not bulldozed to fit a predetermined model.

Extends hardware life cycles, reduces e-waste and manufacturing emissions

Significantly lower upfront cost than new equipment

Aligns with circular economy and Indigenous land stewardship values

Performance ceiling is real but it's often acceptable for community-scale ML workloads

Vendors and networks exist. This isn't a DIY proposition.

Tradeoffs Every Community Needs to Understand

There's no universally right answer.

Every infrastructure pathway involves tradeoffs.

This research makes them visible so communities can decide for themselves.

"Infrastructure ownership must be paired with data governance. Ownership alone is not enough."

– Sundararasan (2026)

Core tradeoff pairs

If you prioritize...	You may be trading away...
Sovereignty (local hosting)	Maintenance burden, technical expertise, higher upfront cost
Affordability (cloud)	Sovereignty, data control → partnership ≠ governance (Wyoming PM, 2026)
Low emissions (small models)	Capability, flexibility → right for language work, limits for complex tasks
Low technical burden (managed services)	Independence, cost over time, long-term accountability
Circular hardware (refurbished)	Performance ceiling → acceptable for most community ML workloads
Cooperative ownership	Organizational complexity, financing, uptime and cybersecurity capacity

What it takes to replicate

What makes it fail

Replication requires:

Sustained funding

Not one-time grants. Operations must outlast launch

Transferable expertise

Non-individual-dependent technical knowledge

Reliable connectivity

Broadband is a prerequisite for any owned compute model

Governance capacity

Institutional structure to own, maintain, and make decisions

Replication fails when:

Funding is one-time and doesn't cover operations (broadband case studies)

Technical expertise is individual-dependent and can't transfer

Connectivity is inadequate. Compute sovereignty requires reliable broadband

Governance capacity can't sustain the system long-term

A Decision Toolkit for Communities

Match your priorities to pathways.

Community priority profile	Best-fit pathway (grounded in research)
High sovereignty + low technical burden	Refurbished local servers + small models (Nubo-style)
High sovereignty + higher technical capacity	Local data center or cooperative compute (platform.coop model)
Affordability first, sovereignty acceptable	Green cloud co-location with sovereignty clauses in contract
Environmental integrity + low capability needs	Edge compute + air cooling + small models (Alipio moku model)
All four priorities (phased approach)	Small models now → cooperative ownership over time

The Haudenosaunee Pod: proof of concept in your own network

What the survey told us

Compute	Community-owned server + Raspberry Pi edge devices + Western University SharePoint
Data control	Full control: The only pod in the network that said this
Sovereignty score	5/5 and infrastructure already matches it
Standardization	Hardware AND software, not just framework. A packaged approach simplifies replication and tackles cost over time.
Stated goal	Refuse any attempt at sovereign ownership from non-Indigenous interests.
Scale philosophy	"Think small-scale. Build local data centers to fit community needs, not the globe's."

A direct ask — to this pod, in this room

You've built what the rest of this network is trying to build. The most valuable thing you can do for Abundant Intelligences right now is document it.

Four questions:

1. What hardware did you use and what did it cost?
2. What expertise was actually required?
3. What's in your standardized package so far?
4. What would you tell the next pod to do differently?

A starting point for every pod

Starting on laptops + hard drives

e.g., T'Karonto, Niitsitapi

→ Run your first language model locally

Install Ollama (free, open-source) on hardware you already own. Test a small multilingual model against your language corpus. Pair the experiment with an OCAP-grounded conversation about what data should never leave the community. No cloud, no new spend, no IT required.

Want to reduce environmental footprint

e.g., Niitsitapi, T'Karonto

→ Calculate your current compute footprint

Use the Green Algorithms tool (green-algorithms.org – free) to estimate your current energy use by workload type. Takes an afternoon, costs nothing. The result gives you a baseline so you can measure whether any infrastructure change actually reduces impact and what to prioritize first.

Using cloud deliberately for experimental work

e.g., Wihihanble S'aat Center

→ Map which workloads actually require cloud

Your dual setup (own server + commercial cloud) is a design choice, not a failure. Make it explicit. Write down which workloads run where and why, with a clear rule for what never goes to commercial cloud. Protects the experimental space and the sovereign space at the same time.

Already have local infrastructure

e.g., Haudenosaunee Pod, Hiringa Te Mahara

→ Document your setup as a replication package

Write up your hardware list, costs, setup steps, and what failed. Even a two-page document makes your model replicable by other pods. The Haudenosaunee Pod is already doing this. Reach out to mziegle2@uwo.ca to collaborate rather than duplicate the work.

The dependency this research didn't fully solve

*"The single largest hidden dependency in Indigenous AI research right now is not cloud compute, **it is the pretrained models themselves**. Even a fully offline pipeline still typically begins by downloading a model trained on corpora the community had no role in shaping."*

— T'Karonto Pod

What this means for the network:

- Infrastructure sovereignty and model sovereignty are two separate problems. Solving the first doesn't solve the second.
- Any pod running an 'offline' pipeline on a locally-hosted server is still dependent on a model trained externally, often in the U.S. or Europe, on data the community didn't consent to.
- This research addressed the infrastructure layer. The model supply chain — who trains the models, on what data, for whose benefit — is the next frontier.

40+ sources across 6 research categories + research pod surveys

8

Data Sovereignty & Governance

Sundararasan 2026; Native BioData; Sargsyan 2022

8

Energy & Power Ownership

Science 2025 Tribal Energy; Energy Mix 2026; CER 2026

3

Community Infrastructure

Stateline 2024; Ford Foundation 2026; Internet Society 2024

11

Data Center Alternatives

Nubo; ODINN; Platform.coop; Precedence Research 2025

7

Water & Cooling

Sustainability Mag 2025; SWAC (Elahee 2013); Iceotope

3

Policy & Partnership Risk

Wyoming PM 2026; BIA TERA; WRI 2026

Responsible AI infrastructure is not one solution.

It's a set of infrastructure tradeoffs communities should be empowered to choose for themselves.

One concrete action

See the 'a starting point' slide. There's one for your pod's situation.

One connection to make

Reach out to the Haudenosaunee Pod before building anything. (mziegle2@uwo.ca)

One open question

Infrastructure sovereignty and model sovereignty are two different problems. This research solved the first.

Research Team

Lead MBA Researchers
Bard MBA in Sustainability



Emerson
Sarmiento



Karla
Pinzón



Lenna
Stahl

Project Sponsor
Abundant Intelligences



Suzanne
Kite



ABUNDANT INTELLIGENCES

Thank you!

Q&A?