


Reclaiming the Stack

Co-operative AI Infrastructure for
Canada's Solidarity Economy:

A Feasibility Study



presented by

HYPHA

CanTrust
HOSTING.COOP



Reclaiming the Stack

Co-operative AI Infrastructure for Canada's Solidarity Economy: A Feasibility Study

Executive Summary	01
Problem Statement	04
Part 1: Target Audience	08
Part 2: Potential Approaches	13
Community-Managed AI Cloud	14
The Sovereignty Suite	15
Case Studies & Comparisons	18
Part 3: Financial Viability & Strategic Business Models	21
Part 4: Technical Feasibility	23
Technical Requirements of AI	23
Part 5: Risk Assessment	31
Conclusions: From Digital Dependency to Co-operative Autonomy	34
AI Disclosure Statement	37
Glossary	37
Appendix A: Energy Use Graphs	39
Bibliography	39





Executive Summary

Every organisation in the solidarity economy is confronting the same uncomfortable question: how do you adopt a technology that could transform your operations when that technology is owned, governed, and monetized by the corporate interests your sector exists to counter? Artificial intelligence (AI), by which we mean generative AI, primarily large language models (LLMs), is no longer optional infrastructure; it is embedded in the productivity tools, communication platforms, and administrative systems that co-operative organisations rely on every day.

This study, produced by CanTrust Hosting Co-operative and Hypha Worker Co-operative with funding from Co-operators' Co-operative Development Program, examines whether a genuinely co-operative alternative is feasible: one that preserves data sovereignty, democratic governance, and environmental integrity without asking organisations to accept inferior tools or prohibitive costs. Our report covers market analysis, technical feasibility (including original research on energy consumption), risk assessment, and financial viability.¹ The pages that follow make the case that such an alternative is both technically achievable, and clearly needed.

The Urgent Reality

The AI landscape is currently dominated by proprietary platforms that conflict with co-operative values. AI assistants are integrated into most of our applications, collecting personal and organisation usage data for a handful of corporations. News articles decry the disruption of numerous professions while the venture capital driven hype machine makes it difficult to understand AI's wider relevance. In late 2025, Hypha and CanTrust surveyed individuals and organisations from the solidarity sector, and this survey data paints a stark picture of the sector's current vulnerability:

- **The Threat of "Shadow AI" is Immediate**
67% of respondents are using AI in the workplace, yet less than a third of surveyed organisations have a formal AI usage policy. This has led to over half of employees using commercial tools without permission, creating invisible data leaks that expose sensitive information to third-party corporate systems.
- **Sovereignty Requires Ownership**
So-called "Canadian" AI services that rely on commercial API wrappers for American models do not provide true sovereignty. Data processed through these applications remains vulnerable to foreign legislation, including the U.S. CLOUD Act and the Foreign Intelligence Act (FISA). Sovereignty compliance demands physically hosting data on infrastructure owned by Canadian-controlled organisations, and shielding sensitive data from foreign surveillance.

Providing a verifiable, secure alternative is no longer just an ethical preference; it is an immediate operational necessity to prevent data leakage and compliance failures.

¹ Note: While we offer operational and technical insights we do not include any formal legal risk analysis.





The Solution: The Canadian Solidarity Stack

To achieve digital autonomy, the co-operative sector must move from consuming AI to owning the intelligence infrastructure that powers it. This report outlines a phased roadmap toward that goal: beginning with immediately actionable alternatives to corporate AI tools (Phase 1), progressing to enterprise-grade options through managed cloud services (Phase 2), and culminating in the full Canadian “Solidarity Stack”² (Phase 3), a long-term, sustainable AI ecosystem built for and governed by Canadian organisations. To reach this goal we need to own each layer of the AI stack, which can be categorised into three distinct, manageable pieces:

1. The Infrastructure Layer (The Community-Managed AI Cloud)

The foundational layer. True data security demands physical hosting on infrastructure owned by Canadian-controlled organisations. This layer involves pooling compute resources on Canadian servers, ideally powered by verified renewable energy, and centrally managed by a co-operative provider to ensure absolute data jurisdiction.

2. The Model Layer (Open-Source & Custom Models)

The middle layer. Instead of relying on proprietary models built on unpaid labour extraction, this layer leverages open-weights and open-source models (such as Llama-3, Olmo, or Apertus). These models can be ethically audited, transparently customized, and run entirely within the Canadian context.

3. The Application Layer (The Sovereignty Suite)

The top, user-facing layer. This suite provides the everyday tools workers need: a secure AI gateway to replace shadow chat tools, an institutional memory engine to query internal co-operative data safely, and a lightweight auditing layer for institutional accountability.

This report grounds the conversation in the reality of our sector through three lenses:

- **Energy Measurement Data**

AI data centers use huge amounts of energy at great environmental cost. Big Tech companies do not disclose this impact. We conducted hands-on hardware testing to measure power consumption in kilowatt-seconds (kW-sec), and evaluated various hardware configurations to identify the optimal balance of operational speed with our sustainability values.

- **Co-operative Sector Survey**

We surveyed organisations across the Canadian solidarity economy, revealing that despite ethical reservations, 67% of respondents are already using AI in the workplace; with over 50% identifying “values alignment” and “security and privacy” as main barriers to organisational adoption. We also conducted one follow-up interview to get a better sense of organisational AI use.

- **A Canadian Focus**

We researched AI deployment in the context of Canadian data sovereignty. Of concern, Canadian data is at risk because of foreign legislation, such as the U.S. CLOUD Act and FISA. Fortunately, Canadians have access to some of the world’s most sustainable and legally secure hosting environments.

² Trebor Scholz, “The Solidarity Stack” (keynote, Cooperative AI Conference, Istanbul, Nov. 11, 2025), Platform Cooperativism Consortium, Dec. 21, 2025, <https://platform.coop/blog/the-solidarity-stack/>. Note: Morshed Mannan is co-credited with coining this term.





The key takeaways for decision-makers include:

Ethics and Security Define the Market

Solidarity economy organisations are adopting AI at five times the national Canadian average (over 60% in our survey versus 12% from Statistics Canada’s June 2025 survey³) despite having ethical reservations about the tools currently available. Of those, 87% cite “values alignment” as their top consideration when selecting AI tools: this is a large, motivated, underserved market looking for an alternative that doesn’t yet exist at scale.

Sustainability and Performance Can Be Balanced

Many people significantly overestimate how much energy an AI query uses: our testing found a basic prompt consumes the equivalent of 1.5 to 6 seconds of microwave cooking. Hardware and model choice are the more meaningful variables. In our testing, Apple Silicon was the clear efficiency leader: a \$1,500 Mac laptop can handle local inference (where the entire process happens on your computer, not the cloud) for individual use, and a \$5,000 Mac Mini Max is a viable production-grade server for shared internal needs. As well, we found that switching a model to “thinking” mode doubles both energy use and run time, so choosing the smallest model for the task is the best way of keeping consumption in check.

Demand Is Demonstrable, With Caveats

Our survey data confirms a latent market: A majority of respondents cite “values alignment” as their top consideration when selecting AI tools, and shadow AI use is a genuine compliance and data security problem that organisations need to solve. However, with only 43% of respondents willing to pay a premium for a values-aligned tool, this market will not be won on ethics alone. The co-operative AI products most likely to succeed are those that lead with affordability, turnkey implementation, and French-language capability, and treat values alignment as a differentiator rather than a price justification.

A Clear Roadmap Exists

The path to co-operative AI infrastructure doesn’t require a leap of faith: it starts with affordable, widely-available hardware that organisations can deploy today. From there, a structured three-phase roadmap moves through shared, managed co-operative cloud hosting toward the full Canadian Solidarity Stack: an ambitious but sovereign, community-governed AI ecosystem built for and by the sector.

In the future, a Canadian’s right to reliable computing, digital literacy, and access to publicly funded AI might be viewed as fundamental as the right to clean air and water. When we speak of “co-operative AI,” we are talking about changing the ownership and governance of the AI infrastructure, relying on an ecosystem of open-source, locally governed, sustainable, and ethically audited digital tools designed to serve communities effectively. However, without immediate investment in our own alternative infrastructure, the solidarity economy risks being left on the wrong side of a technological divide.

We call upon Canadian co-operative leaders, federations, and sector infrastructure organisations to actively fund and support Phase 3 of this roadmap. By pooling our resources to build the Community-Managed AI Cloud, we can collectively move from digital dependency to true co-operative autonomy. A values-aligned AI ecosystem is not only technically feasible: it is urgently required.

3 Statistics Canada, “Workplace adoption of AI in Canada,” June 2025, <https://www150.statcan.gc.ca/n1/pub/11-621-m/11-621-m2025008-eng.htm>.



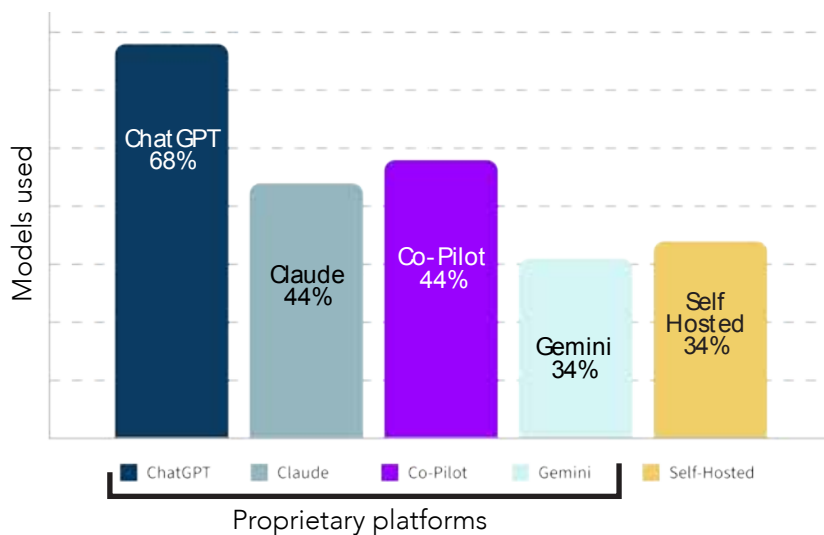


Problem Statement

Organisations in the solidarity economy, which includes the co-operative sector, face a growing, complex set of barriers to adopting artificial intelligence tools that align with their values and operational needs. These challenges are in many cases compounded by macroeconomic trends that have seen budgets reduced while demand for services grows. In this context, administrative efficiencies and automation of routine digital tasks could be immensely helpful; yet existing automation and LLM-based tools are falling short with regards to data privacy, costs, and energy use.

The AI landscape is currently dominated by a handful of large, proprietary platforms (e.g. ChatGPT, Claude, Copilot, Gemini) developed outside the co-operative and social economy context. These tools are developed primarily for profit-driven enterprises and are often built upon what respondents to our November 2025 survey identified as the “unpaid extraction of labour, copyright infringement, and a centralized political project.”⁴ Few alternatives exist that are designed with the principles of democratic governance, collective ownership, or community benefit in mind. As a result, organisations in this sector choose between foregoing AI tools altogether or adopting systems that conflict with their foundational values.

Models Used at Work



Which models are you using? Commercial frontier models dominate current use, reflecting the absence of accessible, values-aligned alternatives — not an endorsement of those tools by the sector.⁵

4 Co-operative AI Blueprint Survey, CanTrust Hosting Co-operative and Hypha Worker Co-operative, November 2025.


5 All graphs and charts use data from the Co-operative AI Blueprint Survey, issued by CanTrust Hosting Co-operative and Hypha Worker Co-operative, November 2025, unless otherwise stated.





The lack of alternatives is compounded by serious concerns about privacy and data security. Many solidarity economy organisations work with vulnerable populations, handling highly sensitive beneficiary data (e.g., asylum applications, health records) that require data jurisdiction. Existing commercial AI platforms offer limited transparency, create ‘API Debt,’ and are prone to surprise updates that can unexpectedly alter how data is processed. That being said, AI tools can help with administrative work, research, and process-related efficiencies. Our survey data reveals a trend of shadow AI use as workers reach for freely available chat-based AI tools.⁶ In organisations without an AI policy, over half of the workers are using commercial AI tools for work regardless, risking data leakage and compliance issues.

AI Usage Policies



Does your organisation have an AI usage policy? Fewer than one in three surveyed organisations has a formal AI usage policy, yet 60% are using AI at work. This gap between use and governance creates risk.

Beyond ethical and privacy concerns, there is an additional operational barrier: the lack of verifiable information in responses. Survey respondents report that the time saved by using LLMs for research and related queries is often negated by the work of fact-checking hallucinations and unreliable data. And while this particular challenge is lessening as models mature and citations improve, general-purpose models still often fail to reflect the nuances, cultural contexts, and specialized methodologies unique to the social impact sector. Without tools that offer confidence scores, transparent citations, and the ability to securely query internal data, the technology fails to meet the accuracy standards required by co-operatives.

Another significant barrier is the environmental impact of AI. The staggering energy consumption, carbon footprint, and water usage of the data centres powering AI conflict with the sustainability mandates of the solidarity economy. For a portion of the sector, the “ridiculous power usage”⁷ of mainstream AI is not just a concern, but a deal breaker that prevents adoption entirely.


6 Michael Rowinski, “Shadow AI: the hidden threat quietly undermining your business,” Mimecast, April 21, 2026, <https://www.mimecast.com/blog/shadow-ai-the-hidden-threat/>.

7 Co-operative AI Blueprint Survey, CanTrust Hosting Co-operative and Hypha Worker Co-operative, November 2025.






Barriers to adopting AI



What are the main barriers to adopting new AI tools in your organization? Values alignment and security and privacy dominate as barriers, cited by over 50% of respondents, outpacing cost and technical capacity for respondents.

Finally, there is a significant technical capacity and infrastructure gap. While local-first LLMs offer a promising, privacy-by-architecture solution for off-grid or highly secure work, co-operatives face hardware bottlenecks, lacking the capital for the expensive GPUs and VRAM required to run these models locally. (This challenge isn't unique to the co-op sector; [chip availability](#) is decreasing globally). Furthermore, staff and board members often lack access to AI literacy education tailored to their context. General-purpose training programs do not address the specific democratic governance questions or co-operative values relevant to this sector. A participatory, co-design approach, driven by membership will likely be more productive than adopting a top-down, managerial style when introducing AI tools.

Data Centre Energy Concerns



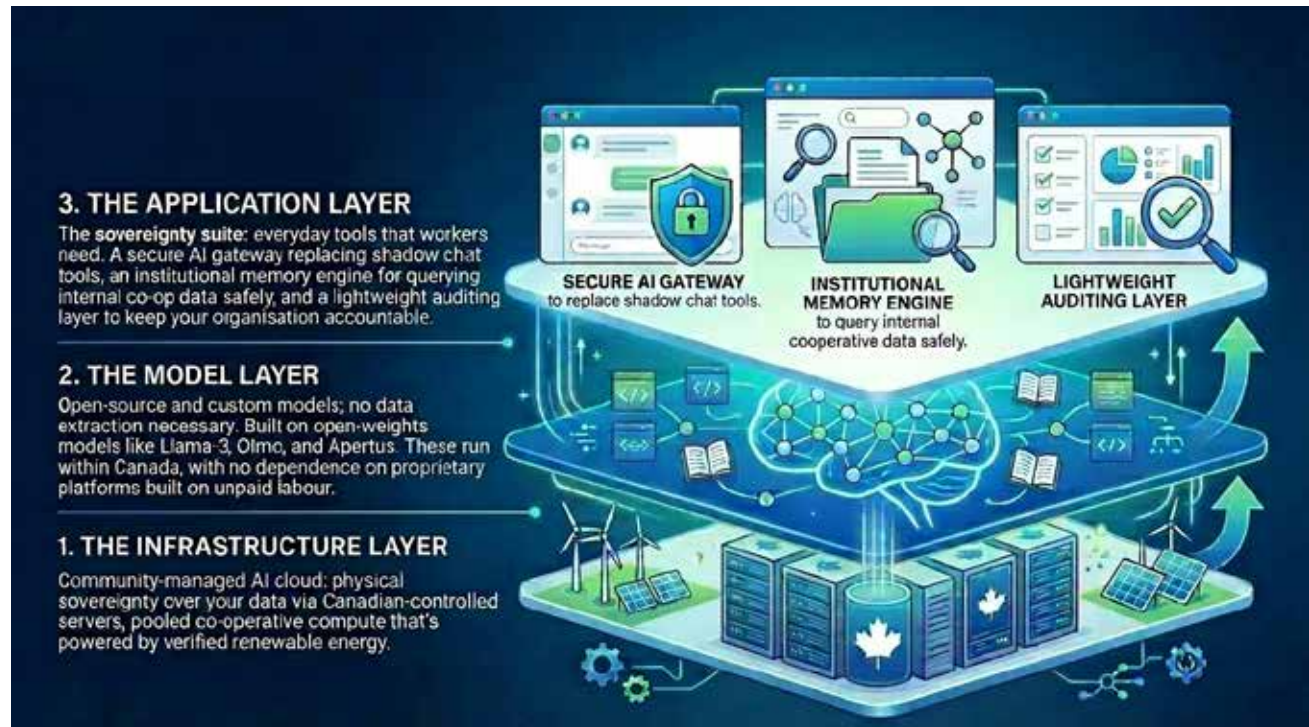
How concerned are you about the increasing energy consumption of AI data centers? A significant share of respondents flagged energy consumption as a serious concern.





Together, these barriers risk deepening the technological divide between the solidarity economy and the broader market. Without investment in an alternative “[Solidarity Stack](#)” of open source, locally governed, energy efficient, and ethically audited intelligence tools, organisations are poorly positioned to evaluate, adopt, or govern AI tools responsibly, ultimately limiting their ability to scale their impact, compete for talent, and serve their communities.


The Canadian Solidarity Stack





Part 1: Target Audience

The solidarity economy is defined as “an alternative model that priorities cooperation, social equity and environmental sustainability over profit and competition. It includes practices like worker co-operatives, fair trade and ethical finance to promote shared prosperity and empower communities.”⁸ Here, we’re focusing on Canadian co-operatives to better gauge how or if an AI Solidarity Stack or elements thereof, might come to exist.



Sources: Statistics Canada, Business-Linkable File Environment, Business Register; Innovation, Science and Economic Development Canada, list of non-financial co-operatives; and national, provincial and territorial co-operative registries.

© His Majesty the King in Right of Canada, as represented by the Minister of Industry, 2026
Catalogue number: 11-627-M | ISBN: 978-0-660-98952-5



www.statcan.gc.ca



8 Population Matters, “Beyond profit: An introduction to the solidarity economy,” November 2024, <https://populationmatters.org/news/2024/11/beyond-profit-an-introduction-to-the-solidarity-economy/>





Canada's co-operative and mutuals sector is a meaningful but frequently overlooked part of the national economy. As of 2024, there were over 6,500 co-operatives and/or mutuals in Canada, with 45% located in Québec due to its enabling legislative environment and mature solidarity ecosystem.⁹ Outside of credit unions and mutuals, there were 5,714 active non-financial co-operatives in Canada in 2023.¹⁰ These organisations contribute \$52.8 billion to Canada's GDP (about 2.5% of the national total). At this scale, the sector is substantial enough to represent a real market, but concentrated enough to approach strategically. Co-operatives are also fairly resilient: while the total number of organisations declined slightly throughout the COVID-19 pandemic, revenues grew 9%, assets grew 22%, and employment rose 7% between 2019 and 2021, all outpacing the broader economy.¹¹

In other words, this is a stable sector, showing some financial growth, which is a reasonable indicator of the capacity to invest in new tools. Moreover, recent research by Co-operatives and Mutuals Canada shows that general support for the co-operative model is high, with one-third of Canadians belonging to a co-op, and another one-third saying they'd like to join one.¹² Canadians also associate co-ops with community-level stability; 81% see co-operatives as part of the solution to affordability and inequality.¹³ Although our study is focused on the feasibility of developing AI products and services within the sector, given the current period of economic and social instability, co-operatively developed technology solutions may be able to capitalise on the general positive perception and goodwill outside the sector as well.

Canada is a country of small and medium sized businesses, and the co-operative sector is no exception. With the exception of 26 organisations, every co-operative and mutual in Canada employs fewer than 500 people, the vast majority have fewer than 99, and nearly half of non-financial co-operatives have no employees at all.¹⁴ Coupled with this, total salaries across the non-financial segment were \$2.7 billion in 2021, which given the spread across thousands of organisations, underscores that most individual co-ops are operating with lean teams and constrained budgets.¹⁵ Québec deserves a note of its own: 45% of all co-operatives and mutuals are located there, reflecting its distinctly enabling policy environment and a stronger solidarity culture.

There are two meaningful segments when thinking about AI fit: the financial segment — credit unions, caisses populaires, insurance co-operatives, and mutuals — is smaller in number but larger in economic weight, with the highest assets, salaries, and the biggest employer footprint. These organisations are more likely to have the budget, regulatory exposure, and operational complexity to adopt AI tools for member services, compliance, fraud detection, or back-office automation. In contrast, the non-financial segment is more numerous and diverse, spanning real estate and housing (the largest group at 33.6% of co-ops), wholesale and retail trade (13.2%), and

9 Co-operatives and Mutuals Canada, "The Economic Impact of Canadian Co-operatives and Mutuals," Nov. 20, 2024,

10 Statistics Canada, "The economic impact of Canadian co-operatives and mutuals, 2021," 2021, <https://www150.statcan.gc.ca/n1/pub/11-627-m/11-627-m2026014-eng.htm>.

11 Co-operatives and Mutuals Canada, "The Economic Impact of Canadian Co-operatives and Mutuals," Nov. 20, 2024, <https://canada.coop/en/the-economic-impact-of-canadian-co-operatives-and-mutuals/>.

12 Co-operatives and Mutuals Canada, "What New National Research Reveals About Canadians and Stability," March 25, 2026, <https://canada.coop/en/what-new-national-research-reveals-about-canadians-and-stability/>.

13 Ibid

14 Co-operatives and Mutuals Canada, "The Economic Impact of Canadian Co-operatives and Mutuals," Nov. 20, 2024, <https://canada.coop/en/the-economic-impact-of-canadian-co-operatives-and-mutuals/>.

15 Statistics Canada, "Canadian Non-Financial Co-operatives, 2021," The Daily, Nov. 6, 2023, <https://www150.statcan.gc.ca/n1/daily-quotidien/231106/dq231106c-eng.htm>.

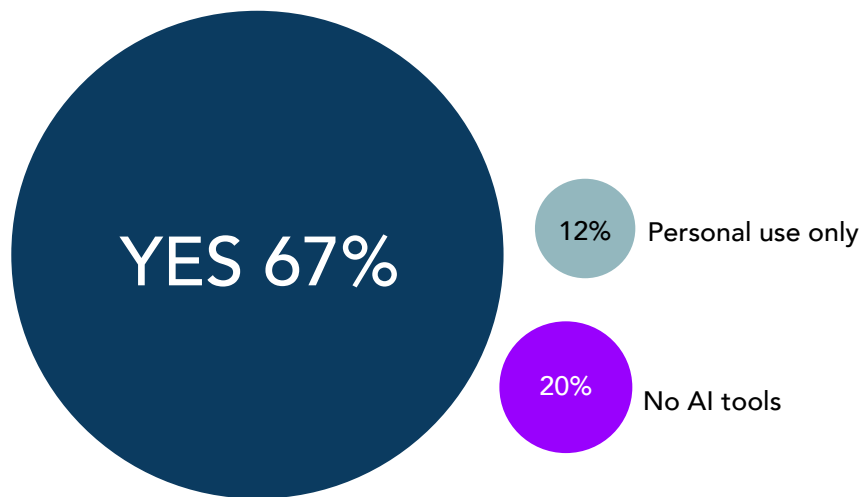




public administration and services (12.6%).¹⁶ These tend to be smaller, more locally focused and are better matched to affordable, general-purpose tools for communication, administration, and member engagement.

A key question for this audience is how they perceive the AI tools and services flooding the technology market. Although overall Canadian AI workplace adoption is pegged at an extremely low 12% (as of June 2025)¹⁷, our survey of solidarity economy organisations and their AI use revealed considerably higher usage with 67% of respondents saying they are using AI in the workplace.

AI Use in the Solidarity Economy



Are you using AI as part of your work? Solidarity economy organisations are adopting AI at nearly five times the national Canadian rate. Over 60% of surveyed organisations reported using AI at work, with the majority using frontier models.

Evidence of higher usage was also presented at the [National Summit on AI and Culture in March 2026](#). While Compétence Culture’s survey of Québec cultural organisations and individual creators reported that “57% do not know how to use AI, do not use it, or are unsure whether it is used within their organisation,”¹⁸ conversely, this indicates that even among those most resistant to AI use (e.g. the cultural sector), upwards of 43% of organisations are using these tools. Thus, it follows that although strong anti- extractive AI sentiments are present within co-operatives and other organisations within the solidarity economy,¹⁹ there is a potential appetite for [value-aligned](#), ethical AI tools and services. Our survey respondents indicated that they are using AI in

16 Statistics Canada, “The economic impact of Canadian co-operatives and mutuals, 2021,” 2021, <https://www150.statcan.gc.ca/n1/en/pub/11-627-m/11-627-m2026014-eng.pdf>.

17 Statistics Canada, “Workplace adoption of AI in Canada,” June 2025, <https://www150.statcan.gc.ca/n1/pub/11-621-m/11-621-m2025008-eng.htm>.

18 Compétence Culture, “L’IA en culture : Mieux comprendre pour agir ensemble,” Nov. 21, 2025, https://competenceculture.ca/wp-content/uploads/sites/2/2025/11/ia-etude-21-novembre-2025_competence_culture.pdf.

19 SOCAN, “SOCAN’s AI ‘No Exceptions’ campaign receives overwhelming support,” n.d., <https://www.socanmagazine.ca/news/socan-ai-no-exceptions-campaign-receives-overwhelming-support/>; Playback, “ACTRA warns CRTC of AI threat to on-screen performers,” May 23, 2025, <https://playbackonline.ca/2025/05/23/actra-warns-crtc-of-ai-threat-to-on-screen-performers/>.





their work for coding and development work, workflow automation, debugging Excel sheets, and summarizing large documents and other tasks.

“ Gemini in Google Sheets has been a game changer for quickly assisting with formulas and chart visualization

“ Summarizing long documents or other pieces of information

“ The coding assistant in Visual Studio is definitely quite neat

Are there elements of your current AI tools you are very satisfied with? Respondents reported genuine satisfaction with AI for time-saving tasks such as coding assistance, document summarization, and workflow automation, but also flagged accuracy and values alignment as ongoing concerns.

Factors When Choosing AI Tools



What factors have influenced or will influence your choice of AI tools? Values alignment leads all other selection criteria at 87%, followed by accuracy of results (79%) and security and privacy (71%).





What might motivate organisations to either switch to a solidarity stack-driven alternative or begin to use AI more frequently? When selecting an AI tool “Alignment with personal/organizational values” rated as the highest concern for our surveyed group (87%), with “Accuracy of results/queries” (79%) and “Security and privacy” (71%) coming next in the list of priorities. For mission-focused organisations, the values and ethics of their digital tools are of increasing importance. Although most co-operatives are likely using Big Tech enterprise solutions (made readily available through programs such as TechSoup), the prevalence of the AI hype, perhaps coupled with discussions of AI’s energy footprint, has elevated values as a product differentiator. This is especially, but not exclusively true when considering the data that has been used to train LLMs; i.e. was it procured in a consensual manner etc. However, overall consumer patterns affect the viability of ‘ethical’ AI tools.

While our survey respondents placed value alignment as a key consideration, only 43% answered in the affirmative to paying more for a value-aligned tool or service, with 27% indicating “Maybe/Depends on the cost.” This is a clear indicator that using values as the key product differentiator is unlikely to translate into sales. AI products aimed at this sector need to also offer turnkey implementation and be consistently affordable. With the exception of the better capitalised financial institutions, enterprise pricing models or solutions needing dedicated IT staff will miss the mark. And given the dominance of Québec, any sector-wide AI product offering must include French-language capability as a baseline requirement.

“94% of respondents consider understanding ethical issues [associated with AI] as the most important and best-mastered AI skill.

- *Compétence Culture, LIA en culture : Mieux comprendre pour agir ensemble*

Apart from the considerations of what to buy, there is also the issue of who would be making purchasing decisions within the target audience organisations. Unlike top-down procurement seen in traditional organisations, in the co-operative sector the separation between AI buyers and users is defined by a unique democratic tension. While buyers, ranging from worker co-ops seeking operational efficiency to consumer co-ops focused on member support, are looking for institutional stability and ROI, users are often mission-driven knowledge workers, grant writers, or developers. This group is currently using shadow systems to keep up with demanding workloads but may still be resistant to adopting Big AI solutions. Unlike a standard SME where a manager might simply mandate a tool, a co-op must bridge the gap between the buyer’s need for a sanctioned, professional solution and the user’s requirement for a safe, values-aligned alternative that protects their collective data.





Part 2: Potential Approaches

The global expenditure on AI is projected to reach ~\$2.52 trillion dollars in 2026²⁰; both a mind-boggling sum for a subset of the overall technology market and a daunting number for anyone considering launching a new product. While much of the hype and focus continues to be on cloud-based proprietary solutions, there has also been a notable increase in discussion about sovereign AI, local or edge models, or simply a return to on-premise hardware. Mozilla Foundation noted that sovereignty was a key theme of the February 2026 India AI Impact Summit, with countries signaling “the same goal: more autonomy over data, more choice in suppliers, and less dependence on a handful of players.”²¹ Respondents to our survey also indicated that data protections and sovereignty were an important consideration in their AI use, while expressing some frustration with existing options. However, in our follow-up discussions with co-operative members, it also became clear that affordability and simplicity are non-negotiable.

Average Satisfaction With AI Tools at Work



KEY TAKEAWAY

Users are cautiously optimistic but remain frustrated by “noise” and hallucination issues.

Average satisfaction with AI tools at work Overall satisfaction is moderate, with affordability and values alignment the most commonly cited gaps between current tools and sector needs.

Given these constraints, and the heterogeneous nature of the co-operative sector - not to mention the solidarity economy taken as a whole - it’s difficult to elucidate a discrete set of products or services. In this section, we’ll outline two approaches to creating sovereign, performant AI for the Canadian solidarity economy: the *Community-Managed AI Cloud* and the *Sovereign Suite*.

20 Business Insider, “Private LLM growth expected as enterprises shift GenAI from experiments to secure domain-specific systems,” 2026, <https://markets.businessinsider.com/news/stocks/private-llm-growth-expected-as-enterprises-shift-genai-from-experiments-to-secure-domain-specific-systems-1035769007>.

21 Mozilla Foundation, “A dispatch from the India AI Summit floor,” Mozilla Foundation Newsletter, March 4, 2026.





Community-Managed AI Cloud

The infrastructure layer of the Solidarity Stack²² starts from a simple premise: Canadian co-operatives shouldn't have to choose between the convenience of cloud AI and their values of privacy, sovereignty, and environmental integrity. The approach involves pooling compute resources across a small number of anchor clients, hosted on Canadian servers running on verified renewable energy, but managed centrally by a co-operative provider on behalf of member organisations. Each co-operative gets a private, segmented environment for their data, without having their data transit or reside in the U.S.


Shared compute is not novel-the academic sector has long used this approach for research. The community-managed model adapts it for small and medium-sized organisations who will likely only require enough compute for workflow automation, knowledge management, and other types of day to day inference, rather than intensive model training. Renewable energy is a core commitment, not an afterthought. Starting early conversations with Community Energy Cooperative Canada and their members to better understand how and when they might be able to meet this type of demand is crucial, with hydro-electric providers providing a practical interim option. Lastly, community governance is what separates this model from just renting co-operative-branded cloud infrastructure. With the right structure, members have a meaningful voice in pricing, data policy, and service priorities, and the organisation remains accountable to the sector it serves rather than to external investors.

Infrastructure and hosting services primarily compete on cost and support, given that other than up-time and capacity, there is little to differentiate products for non-technical clients. But in Canada (and Europe), there has been a move toward increasing digital sovereignty; in particular, to develop alternatives to the offerings of American companies. As Canadian data sovereignty concerns become board-level issues, not just IT considerations, more co-operatives and nonprofit organisations will be assessing the risks involved with continuing their dependence on U.S. technology, especially related to AI. This may become especially important should the Canadian government put forward legislation to replace the Personal Information Protection and Electronic Documents Act (PIPEDA), which is expected to address data sovereignty. Taken together, the hope is that consumer sentiments regarding digital infrastructure will become more nuanced, with data sovereignty and energy source counting alongside price in purchasing decisions.

A Community-Managed AI Cloud is a long-term vision for a more local and private digital infrastructure. As an interim step, co-operatives and social purpose organisations might consider developing and offering data audit services, helping organisations understand their digital risks with regards to AI. Similarly, for the many organisations who lack internal IT support, affordable fractional CTO (Chief Technology Officer) services offered by organisations and individuals with lived experience in the solidarity economy are an essential stepping stone to greater digital resilience. Imagining that some of this movement-building and education work has already been done, the following chart offers potential starting points for the future Community-Managed Cloud service.

22 Platform Cooperativism Consortium, "The Solidarity Stack," Dec. 21, 2025, <https://platform.coop/blog/the-solidarity-stack/>.





Audience	Price Point (estimated)
<p>Best immediate fit: Mid-size credit unions and caisses populaires, larger co-operatives</p> <p>These organisations have the budget, the regulatory pressure, and the capacity to make a hosting decision on sovereignty and compliance grounds alone. Ideally, they are already asking their current vendors hard questions about data residency. A Canadian-owned, co-operatively-aligned alternative gives their compliance officers something to point to. These organisations are large enough that the economics work: they generate enough AI workload to justify dedicated compute allocation. This is the founding member cohort.</p>	<p>\$3,000-\$6,000/month</p> <p>This covers a fully segmented, dedicated compute environment, managed security and updates, Canadian data residency guarantee, and a defined Service Level Agreement (SLA).</p> <p>At this price point 8-10 anchor clients are required to cover infrastructure costs before the business is viable.</p>
<p>Strong fit with longer cycle: Co-op federations and sector infrastructure organisations</p> <p>Umbrella organisations for the sector, provincial co-op associations, or credit union centrals may see the value in procuring Community-Managed Cloud as a shared service for their members. Given the number of stakeholders involved, this will likely be a slower rollout, creating a solid second cohort of users.</p>	<p>\$800-\$1,500/month</p> <p>Multi-tenant (i.e. shared) infrastructure with strong logical separation, same data residency and environmental guarantees, but with lower compute allocation.</p>
<p>Values fit, budget uncertainty: Worker co-ops in professional services</p> <p>Legal, accounting, design, and consulting worker co-ops care about client confidentiality and might resonate with the sovereignty argument. The challenge is budget: a 12-person worker co-op can't pay for dedicated GPU compute, so the angle is community access (i.e. shared) with less compute allocated than the above tier.</p>	<p>\$150-\$400/month</p> <p>Essentially a values-aligned SaaS subscription. Given the thin margins, this may need to wait until there are enough anchor organisations signed up.</p>

The Sovereignty Suite

This is the application layer of the Solidarity Stack, consisting of an integrated product built from three interlocking pieces: a secure AI gateway that replaces shadow tool use, an institutional memory engine that makes the co-operative's own knowledge queryable, and a lightweight audit layer that provides visibility and accountability over how AI is being used within the organization. These three components are designed to work together, with the gateway (most likely a chatbot) being what members or management use every day. The Suite runs on any compliant





Canadian hosting environment and doesn't require the Community-Managed Cloud (though that combination represents a more fulsome expression of the Solidarity Stack). Organisations with existing robust infrastructure - for example, a credit union with an established private cloud, can deploy the Suite on top of what they already have.

The value proposition is two-fold: first, it eliminates shadow AI use by offering a tool that's as easy to use as ChatGPT, while also providing visibility into how AI is being used internally. Ideally, it also offers access to a variety of open-weights models so users can select one best suited to their task. For many organisations, this resolves an internal tension between productivity and commitment to non-extractive values. Secondly, the Memory Engine component, most likely a RAG-based configuration, provides intelligence grounded in the organization's own data. Institutional memory loss - through worker turnover or poor record keeping - is a challenge across the sector, and the Sovereignty Suite addresses this issue.²³

AI literacy, capacity building and workflow improvements are services that could accompany the adoption of the Suite. Technology-related change is hard in any organisation and industry, and there are persistent misconceptions about AI's capabilities and utilities. Workers and members require support and strategies on how to adopt and use the Suite for specific tasks such as grant writing or membership management. Given that the proposed tool isn't an AGI-type solution, ideally there would be numerous versions of the Sovereignty Suite, focusing on different verticals within the co-operatives sector. Consumer-facing co-ops such as grocery stores, manufacturing co-ops, and agricultural co-operatives have different pain points that would be best addressed with bespoke versions of the Sovereignty Suite.

Like the Community-Managed Cloud, the Suite is an ambitious notion that will face numerous challenges; not the least of which is the crowded AI solutions market and the sheer noise generated by competing options. As successful businesses, co-operatives will assess price alongside any values-related messaging, making cost a significant factor in the design and deployment of the Suite. Lastly, the governance structure of many co-operatives and non-profit organization is such that adoption decisions may require board approval, slowing the sales cycle down considerably and making it difficult for service providers to have certainty about potential adoption. Conversely, the sector has networks that accelerate trust-based adoption, and if co-operatives are talking to each other, especially within provincial associations and federations, a single successful deployment at a visible credit union or well-regarded worker co-op may improve the chances of wider adoption and success. The following section suggests configurations that might work for a Sovereign Suite product.

23 One Hundred Nights, "Institutional memory and AI," n.d., <https://onehundrednights.com/article/institutional-memory-ai/>.





Audience	Price Point
<p>Professional services worker co-ops and mid-size non-financial co-operatives</p> <p>Legal, accounting, communications, and consulting worker co-ops are the best fit for the Suite as a standalone product. Although this is a small segment in Canada, they have confidentiality requirements that make shadow AI use a specific legal risk, not just a values-related concern. These are knowledge-based organisations, so their institutional memory has real value.</p> <p>Housing co-ops and retail food co-ops are also good candidates for the Suite. These organisations are likely dealing with increasing administrative complexity without growing their back-office headcount.</p>	<p>Gateway Only: Entry Tier</p> <p>\$99-\$149/month flat, for organisations up to 25 users. This covers the secure AI interface, basic governance controls, Canadian data residency, and usage reporting, but no document indexing. This price point aims to compete with what staff are currently spending informally on individual ChatGPT/Copilot/Gemini subscriptions.</p> <p>Full Suite: Standard Tier</p> <p>\$299-\$499/month flat, for organisations up to 50 users. Gateway plus the RAG-based Memory Engine with document indexing, the quarterly values alignment report, and onboarding support.</p>
<p>Co-op development organisations and federations</p> <p>These organisations may be equally likely to buy the Suite for themselves, as they are to purchase it on behalf of the organisations they support, so bulk procurement is worth exploring here.</p>	<p>Federation / Multi-Org Tier</p> <p>Custom pricing, structured as an annual contract, likely in the \$2,000-\$5,000/month range covering a defined number of member organisations.</p>
<p>Credit unions and caisses populaires</p> <p>The compliance argument for the Gateway is strongest here, and there is likely an available budget. However, procurement processes, IT governance layers etc. might extend the sales cycle significantly. As well, a fully bilingual application and service delivery is essential to meet the requirements of the Québec market.</p>	<p>Full Suite: Growth Tier</p> <p>\$799-\$1,200/month flat, for larger organisations up to 150 users, with expanded document storage, priority support, and a dedicated onboarding session. This is the tier for credit unions, housing co-ops with significant staff, and larger agricultural co-operatives.</p>





Case Studies & Comparisons

In this section, we'll explore three case studies that offer comparison points for the Community-Managed AI cloud, and suggest tools and approaches for the Sovereign Suite.

Credit Union Service organisations (CUSO)

Canada has a robust vertical of organisations providing technical services to their credit union partners. Although the term 'CUSO' is more widely used in the U.S., we're adopting it here to refer to any credit union-related services. These organisations, often referred to as Centrals, have existed for decades to allow credit unions to focus on financial and client-related services. Credit unions of all sizes, including smaller organisations, are more competitive because of these shared resources.

[League Data](#) is one example of a Canadian tech-focused CUSO. Based in Nova Scotia, it is a co-operative that provides a common core banking platform, cybersecurity, and managed data services for credit unions across the Atlantic provinces. These services ensure that a small credit union in rural PEI has the same tech power as a larger city branch. This model of layered co-operatives presents an interesting corollary for the Community-Managed AI Cloud as none of the existing CUSOs offer AI related services and in recent years have moved to outsourcing some of their digital service offerings to focus on the "business of payments, clearing and settlement, and treasury."²⁴ Framing the Community Cloud as a 'co-operative CUSO for AI infrastructure' connects a new service offering to an established governance pattern that credit union directors and members already know and trust.

Commons Cloud/SomNúvol

This Catalan initiative is an example of cloud infrastructure built for the solidarity economy, offering members access to business applications (including NextCloud, which is an open-source office software suite) and email services. Commons Cloud/SomNúvol started in 2018 with just a handful of members. By the end of 2025, the network had about 203 members: 99 organisations and 104 individuals, of whom more than half are co-operatives and 30% are associations, all part of the broader Catalan social and solidarity economy ecosystem.²⁵

The initiative is built on the principles of community self-management and ecological and economic sustainability, with individual users and organisations becoming consumer members of the multistakeholder co-operative femProcomuns, which holds agreements with other organisations to produce the service. Commons Cloud is differentiated from other FLOSS (Free/Libre Open Source Software) hosting providers by also offering specialised Enterprise Resource Planning and Customer Relationship Management software to its co-operative members. This intercooperation agreement structure, in which allied organisations contribute to production and governance rather than simply consuming a service, is a governance model worth considering for the Canadian context. There are likely many lessons to be gathered from the Commons Cloud example on balancing technical service provisions with co-operative values, and on the time required to develop similar technical 'ecosystems.'

24 Sheila Vokey, "Reflections on a pivotal year for Central 1," Central 1, Jan. 13, 2026, https://www.central1.com/in_the_news/reflections-on-a-pivotal-year/.

25 Mònica Garriga Miret, David Gómez Fontanills, Xavier Martínez Serrano and Enric Senabre Hidalgo, "Social and solidarity economy and shared knowledge: Facilitating and integrating the collaborative adoption of digital commons in Catalan coops," Internet Policy Review, Feb. 6, 2026, <https://policyreview.info/articles/analysis/digital-commons-in-Catalan-coops>.





National AI Research Resource (NAIRR) and the Pan-Canadian AI Compute Environment (PAICE)

Initially established as a pilot in 2024, NAIRR has become a proven national research infrastructure, supporting more than 600 research projects and 6,000 students across all 50 U.S. states. (NAIRR is akin to the Pan-Canadian AI Compute Environment (PAICE), but the latter has [launched more recently](#)). Like CUSOs, this initiative provides a more equitable distribution of technical services to smaller, rural and/or underserved institutions. A parallel can be drawn between these organisations and small and mid-sized Canadian co-operatives; the same structural disadvantage that limits rural universities applies to worker co-ops and small credit unions trying to adopt AI responsibly. Both NAIRR and PAICE are research-focused and therefore likely have more intense compute requirements than what is being envisioned for the Community-Managed AI Cloud. However, these programs also show that a functioning model of shared AI infrastructure demonstrably lowers barriers for under-resourced organisations.

Sovereign Suite Comparisons

Given the current pace of AI product releases, potential competitors for solidarity economy-focused tools and services, such as those outlined in the above Sovereign Suite section, will rapidly emerge and just as quickly be acquired or go out of business. But regardless of sector, the primary competitor for any AI product, tool or service will be one of the Big AI companies, which offer access to frontier models. This includes OpenAI's ChatGPT, Anthropic's Claude, Google's Gemini, and CoPilot for organisations using the Microsoft suite of tools. In particular, the affordances offered by [Open AI's custom GPTs](#) provide an easy path to adding internal documents to a knowledge base for querying. Likewise, for organisations using Google Workspace or Microsoft 365, Gemini and CoPilot's workspace integration provide a direct way to access and query internal documents. Gemini and CoPilot's data and privacy are governed by existing enterprise-level user agreements, making them a likely first choice for many larger organisations. All of the Big AI companies benefit from enormous marketing budgets, high-profile partnerships, and unparalleled financial investment; all things considered, they are formidable opponents in a crowded field.

However, if we assume our target audience is privacy, security, and ethically-minded, then we can shift to assessing smaller, niche products that may be more appealing to mission-forward organisations. Sovereign AI applications are increasing, offering secure, private tools for teams and individuals alike. While options such as [Open WebUI](#) and [LibreChat](#) are focused on a more technically literate audience, [BionicGPT](#), [Msty](#) and [AnythingLLM](#) all aim to be user-friendly and turn-key. Most of these options can be configured for team or individual use, and offer cloud or self-hosted options. While each of the aforementioned tools are quite similar, there is differentiation in the details. Msty features customizations for law, education, and science, while Bionic-GPT is the most compliance-forward, with its website noting that the product meets the requirements for SOC 2, the ISOs for cloud and information security, and is GDPR compliant. All of these products claim to offer similarly performant LLM inference to the frontier models, plus privacy and security, at a much cheaper price point than ChatGPT or Claude.





Platform	Best For	Key Features	Deployment Options
AnythingLLM	Business users, RAG-heavy workflows	Local-first, agentic, workspace-centric	Desktop (Free), Docker (Free), Cloud (\$50/mo)
BionicGPT	Enterprise governance	Compliance-focused architecture; strict audit logs, team permissions	On-premise, Private Cloud
LibreChat	Multi-provider team collaboration	Code interpreter, artifacts (Mermaid, React), MCP support	Web-based, self-hosted
Msty	Secure, localized team workflows	Persona switching, local machine tool connections, MCP support	Desktop, Web-based
Open WebUI	Tech-savvy teams, local model enthusiasts	Extensive plugin system, RAG flexibility, rapid feature releases	Docker, self-hosted

Another approach to sovereignty is to go local. Local LLMs encompass the open-weights models themselves (e.g. Meta’s Llama family, Mistral, and DeepSeek) and the inference engines used to run them on local hardware - that is, workers’ desktop computers or even phones. This product segment is driven by organisations who are interested in data privacy and/or keeping their AI [token costs low](#).

Lastly, there is a related product segment for technical service providers who want to take advantage of open source tools to create knowledge bases from member documents, build intake and triage agents for community services, or streamline administrative work so organisations can focus on service delivery or client relationships. Imagine a co-operative bundling [Dify.ai](#) deployment (a low-code AI app builder) with knowledge base structuring and ongoing stewardship, at an accessible price point. These services are needed if the solidarity economy sector is to evolve under its own terms.





Part 3: Financial Viability & Strategic Business Models

Access to capital and financing is a long-standing challenge in the solidarity economy, and in the co-operative sector in particular. The issues are especially acute for new organisations who are attempting to grow businesses in uncertain economic conditions. Lack of awareness of the co-operative model hampers investment from traditional financial backers, and there are a limited number of dedicated funds to draw from that provide social impact or purpose-driven investing in Canada.

“Many co-op leaders report that they need additional, outside capital and financing to invest in growth strategies; address ongoing operational issues, such as aging infrastructure; and/or tackle modernization costs, such as the rising costs of adopting new or upgrading technology.
- Canadian Centre for the Study of Co-operatives; *The 2024 Top Co-op Issues Survey Report* ²⁶

Therefore, this section suggests potential economic strategies for sustaining the Solidarity Stack, from initial capital investment to operational self-sufficiency. The idea is to decouple product development from immediate profit requirements in order to build long-term, stable AI infrastructure for the co-operative sector. Ideally, this will lessen the requirements for ongoing capital investments, which are challenging for the sector as a whole.

Technical Mutual Aid

As outlined in the Community-Managed AI Cloud section above, a strategy of securing a few initial anchor clients, such as credit unions or larger co-operatives, will both provide some initial investment and inertia for developing the infrastructure. This also presents an opportunity for a cross-subsidization framework that can lead to movement building. For example, tiered membership dues might mean that large credit unions or worker co-ops pay \$40,000+/year with smaller co-ops paying \$500-\$1,500; similar to existing dues structures for many umbrella organisations. In this scenario, excess compute capacity on the co-operatively-hosted infrastructure would be shared across consortium members, so that smaller organisations can benefit from top-tier technology at an affordable cost. Anchor clients will benefit from meeting data residency or AI transparency requirements, by being seen as an innovative first-mover, and from the fixed cost infrastructure.

Capitalisation Requirements

The infrastructure and application layers require different capitalisation strategies. The Community-Managed AI Cloud is the larger proposition, and is well-positioned to draw on the [Canadian Sovereign AI Compute Strategy](#), a \$2 billion federal initiative announced in 2024. This larger proposition requires a coalition to move ahead, with the budget dependent on the

²⁶ Canadian Centre for the Study of Co-operatives, “The 2024 Top Co-op Issues Survey Report,” 2024, <https://usaskstudies.coop/documents/research-reports/2024-top-co-op-issues-survey-report.pdf>.





size and technology requirements of the participants. The Sovereignty Suite, as an application layer, is somewhat simpler to outline, with an estimated initial \$400,000 ²⁷ investment required to develop a production-ready application. This estimate varies considerably depending on the complexity and regulatory requirements of the tool in development; if we use [Rook](#), Hypha’s agentic knowledge management tool as a guide, it has taken around 400 hours to produce a relatively competent tool in closed beta. This initial investment will fund the first 12-18 months of development, integrations, and security auditing. Organisations should explore the available options (e.g. the Canadian Cooperative Investment Fund and others) to understand if these entities might be interested in providing more patient capital to usher in technology built with human values, not extraction, in mind.

Category	Estimated Cost	Objective
Core Infrastructure Development	\$250,000	Ideation, development, and hardening the application to enterprise standards.
Security & Compliance	\$75,000	Third-party penetration testing and data sovereignty certifications (e.g. SOC 2)
Implementation & Support	\$75,000	Developing onboarding documentation and ongoing support as needed via a help desk.

Hybrid Revenue Strategy: Stability vs. Innovation

To avoid the code rot often associated with grant-funded open-source projects, balancing recurring maintenance fees with high-value technical consultancy is a suggested strategy. Monthly subscription fees from the Sovereignty Suite could provide the funds necessary for ongoing server costs, security patches, and basic maintenance. Like most technology, AI-powered workflows require ongoing maintenance; in fact, the current speed of model development makes this a necessity as large language models update, and workflows break when APIs change etc. As long as the client organisations allocate funds to support this work, maintenance can provide a form of stable revenue for technical service provider co-operatives. However, to ensure the AI tools remain competitive as well as functional, technical consultancy work can act as an R&D ‘engine,’ by which implementation fees, custom integration projects, and specialised training serve as a vehicle for innovation. Custom features funded by individual clients would be designed to be merged back into the core Sovereign Suite, benefiting the entire community. Upfront capital challenges notwithstanding, a combination of product and consultancy models will likely offer the most stable path forward.

²⁷ Note: Based on a \$200/hr billing rate, the \$400,000 budget represents approximately 2,000 billable hours, or one year of work for a senior developer.





Part 4: Technical Feasibility

In this section, we evaluate the practical requirements for Canadian organisations to host and maintain AI models that are powerful enough to be useful, but also sovereign, secure, and sustainable. We examine the current AI ecosystem, and assess hardware, software, energy consumption and data management requirements. By framing our approach to technical feasibility through the lens of the co-operative sector, we seek to answer the question, “Can we offer AI solutions that align with our values?”

Technical Requirements of AI

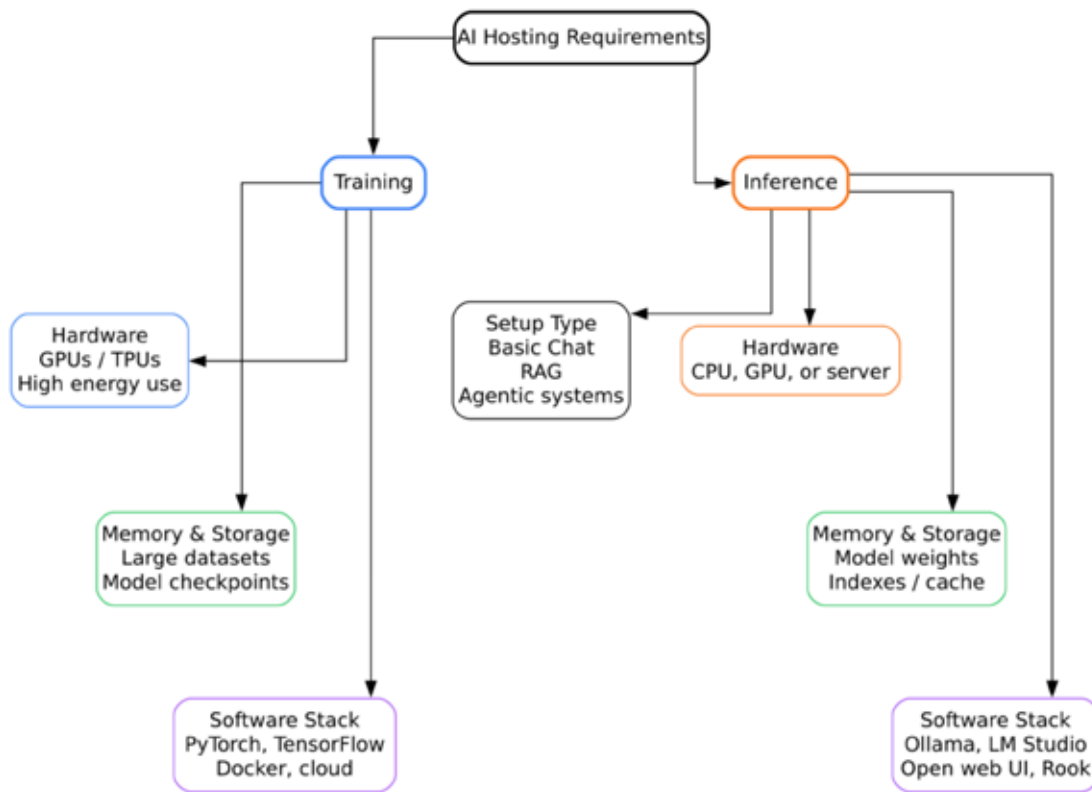
Hosting an LLM requires a combination of hardware and software components, which at a minimum, includes:

- **Computing Hardware:** High-performance GPUs (or TPUs) are typically required for training LLMs. Training is computationally expensive and requires massive energy to power processors. For inference (user queries) modest hardware such as web servers or even laptop CPUs can suffice, depending on the model’s size and complexity. Our focus here is on inference, as the co-operative sector is better positioned to deliver these services, rather than the highly capital- and compute- intensive training of large language models.
- **Memory and Storage:** AI models require significant RAM and storage for both training and inference. For example, models like Llama 3 or Mistral require several gigabytes of memory, with larger models needing terabytes of storage to run.
- **Software Stack:** The software stack for training and development of AI models includes the AI framework (e.g., PyTorch, TensorFlow) and the infrastructure to deploy and manage the framework (e.g., Docker, cloud platforms, etc.). For inference, the software stack requires an LLM runner (e.g., Ollama, LMStudio, Jan AI), the model itself (e.g., Llama, Mistral, etc.), and again the software infrastructure to deploy and manage the model (e.g., Docker, cloud platforms, local operating system, etc.).

The type of AI setup also influences the infrastructure requirements. For instance:

- **Basic Call/Response Chat:** This setup requires minimal hardware and is suitable for simple tasks such as answering FAQs or providing basic customer support.
- **RAG (Retrieval Augmented Generation):** Combining a pre-trained model with a retrieval system to enhance the accuracy and relevance of responses, RAG also requires additional infrastructure for data indexing and retrieval.
- **Agentic Systems:** These are more complex and require significant computational resources, as they involve continuous interaction, decision-making, and learning from user inputs.





AI hosting requirements split sharply depending on what you're actually doing. Training a model demands GPUs or TPUs, massive storage, and high energy use. Running inference (the day-to-day use case for most co-operatives) can be done on a CPU, GPU, or modest server, with significantly lighter infrastructure. This distinction is what makes a co-operative-hosted AI stack feasible.



What do you think consumes more energy: training an AI model once, or the energy used by millions of people using it daily? Respondents' answers reveal a slight misconception about where AI's energy footprint actually falls; a misconception our hardware testing directly addresses.

Power & Energy Considerations

The power required for AI training and inference varies widely depending on the model size and hardware, but frontier AI companies release little to no data about their energy consumption. Organisations that are interested in more ethical and energy-efficient AI options, including CanTrust and Hypha, are exploring ways to minimize the environmental impact of using AI by using local energy-efficient hardware, renewable energy sources, and smaller models, which reduces computational demands. Open source AI hub [Hugging Face](#) has also undertaken similar measurement work.



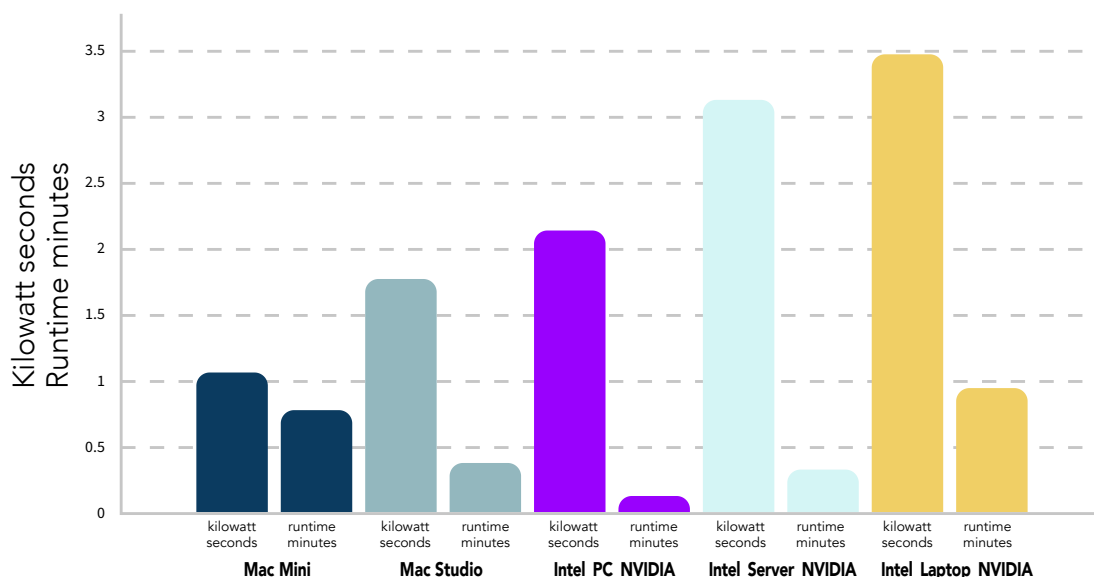


For this report, we performed a series of experiments to gather real data about AI's energy consumption for inference when used locally. We measured power consumption in kilowatt-seconds (kW-sec) to identify an optimal configuration that balances performance with sustainability.

Hardware Comparison for Energy & Speed

The power measurement data illustrates the energy cost per query across different hardware configurations, showing the efficiency of different architectures. What we found was a trade-off depending on the organisational goals: speed per token, power efficiency and scalability are all related. The most efficient hardware choice for an organisation depends on their priorities in this regard.

Energy Use and Run-Time Comparisons



Energy use and run-time for generating a 1,000-word essay. Hands-on hardware testing across multiple configurations shows that energy cost and speed involve a real trade-off - but not an insurmountable one. Apple Silicon leads on efficiency; NVIDIA gaming GPUs lead on throughput. The right choice depends on an organization's priorities.

Maximum Throughput and Performance: Our data shows that high-performance NVIDIA gaming GPUs (specifically those with a 256-bit memory bus - the wider the bus, the faster data can be processed though this is dependent on a few other factors as well) provide the best balance of speed and power for users who cannot wait for responses and want local AI to be as quick as possible. Inference can be acceptably fast even with older GPU generations, with the memory bus width and memory speed being the most important factor. These older generation cards (e.g. NVIDIA 3000 or 4000 series cards) are more readily available and affordable than current generation hardware.

In thinking through configurations, we also need to keep an eye on new, potentially disruptive approaches to model hosting for inference. The Canadian company [Taalas](#) is manufacturing the model layer into custom silicon processing units with mind blowing performance results; orders of magnitude faster than software models.





Efficiency Leaders: Apple Silicon (Mac hardware) is a clear winner for power efficiency, and often readily available if a high-end video card is too expensive or PC hardware unavailable. For individual users, local inference on a \$1500 Mac laptop or Mini can suffice for basic tasks using smaller models, or for prototyping larger systems. A 16GB model is essential for any serious use; the 8GB models are severely limited in what they are able to run. A \$5000 Mac Mini Max meanwhile can act as a feasible production-grade server for shared internal needs, running multiple models at the same time, while retaining a simple user interface (e.g. Ollama, LMStudio, Unsloth Studio, etc).

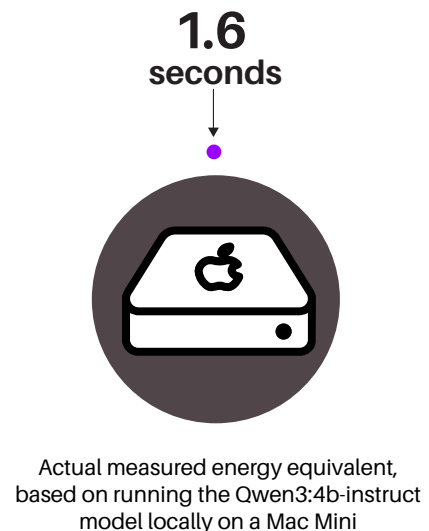
Beyond the Hobbyist - Production Use: While building your own infrastructure at home is a great starting point, making these tools 'highly available' for critical infrastructure requires a more robust approach. In order to cost-effectively serve multiple end-users at once, techniques such as batching and parallelism are used by enterprise providers. It's possible to DIY this in the cloud, but it is challenging to scale the service, especially if more than one model is required. There is also the question of high availability, that is, keeping a model pre-loaded in RAM so it responds immediately, rather than fetching it from storage each time a request comes in. The trade-off is real: a model sitting in RAM is consuming hardware resources around the clock, whether it is actively being used or not.

Energy use for Simple Inference

The Perception



The Reality



Power usage for generating 5 broccoli casserole recipes. We asked survey respondents for their perception of energy use for creating the recipes, and then measured the power usage. Perceived power use was far higher than the actual power usage for a simple query. Depending on the inference hardware, it requires 1.5 to 6 seconds of cooking time to generate 5 recipes.

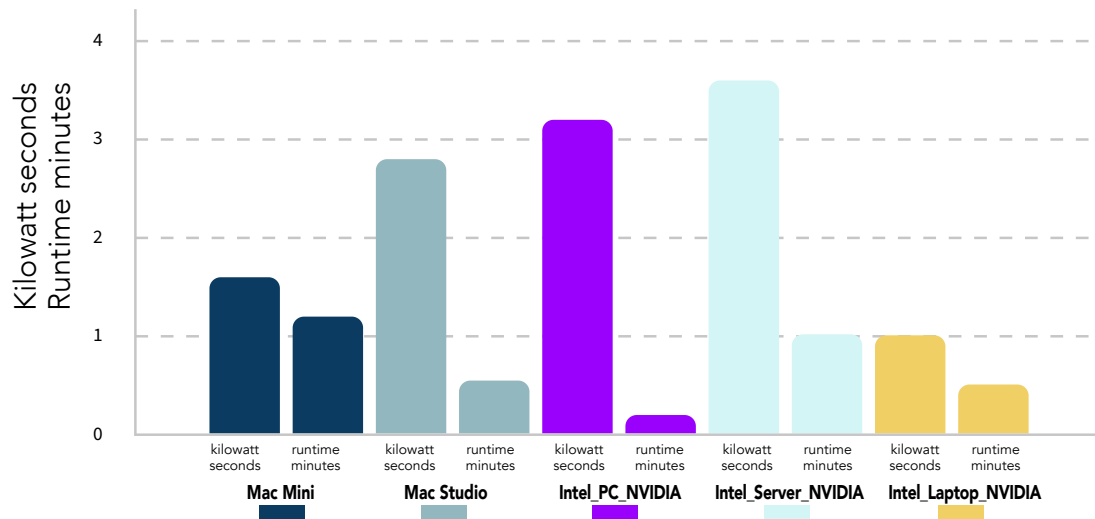




Power Usage and Run-Time for Different Model Configurations

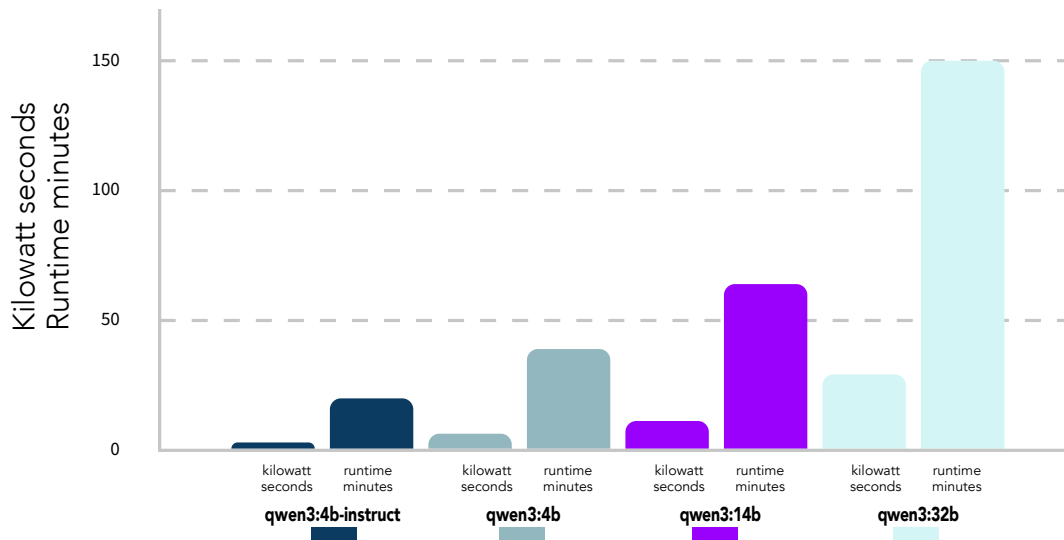
We measured power usage for the same prompt across a range of model sizes, as well as thinking versus non-thinking modes. We found that switching a model to 'thinking' mode doubles both energy use and run time; larger model sizes increase both further. Choosing the smallest model appropriate for the task is a very effective way for organisations to manage their AI energy footprint.

Performance and Energy Use Comparison



Power usage and run-time when generating 5 broccoli casserole recipes across hardware configurations. A simple, everyday AI query such as generating five recipe variations, uses between approximately 1 and 3.6 kilowatt-seconds depending on hardware, the equivalent of a few seconds of microwave cooking. The Intel PC with NVIDIA GPU completes the task fastest (under 10 seconds), while Apple Silicon machines (Mac Mini and Mac Studio) offer the best balance of reasonable speed and lower energy use. Intel's server-grade NVIDIA hardware draws the most power despite not being the quickest. For organizations running routine inference tasks, this data suggests that widely available consumer hardware, particularly Apple Silicon, delivers good performance without the energy overhead of dedicated server infrastructure.





Power usage and run-time across Qwen3 model sizes and modes. Switching the same model from standard to thinking mode (qwen3:4b-instruct versus qwen3:4b, respectively) roughly doubles both energy use and run time. Scaling up model size amplifies this further: the 32b model in thinking mode, where it reasons through problems before responding, uses approximately 14 times the energy and takes 8 times as long as the 4b-instruct (i.e. the standard mode).

Power usage and run time double going from non-thinking (-instruct or standard) to thinking mode, and then both increase with the model parameter size. For most co-ops, choosing the smallest sized model for the job at hand will save the most power and time.

Energy & Environmental Impact

AI is known for its high energy consumption, with an estimated “4.4% of all the energy in the US” now used to power data centres.²⁸ Although the frontier model companies release little information about their energy use, training large language models requires a significant amount of energy for a limited duration, while inference uses more over the long term. Data centres required for AI are also driving conversations about [reopening coal-fired power plants](#), many of which were being phased out in the push toward cleaner energy. And in addition to energy use, there is an active public discourse about data centres and water usage, with competing interpretations of how much water they actually use for cooling (a definitive number is hard to find, but they certainly use some water). High-performance computing hardware generates significant heat, requiring cooling systems that further increase energy consumption; essentially, this is air conditioning for the computer hardware. In locations like the U.S., AI has been identified as a [massive water consumer](#), further exacerbating their clean water challenges. This raises important environmental concerns for organisations committed to sustainability.

Depending on where the AI hardware is located, power sources in Canada can vary from clean (e.g. hydroelectric, renewables, nuclear) to dirty (e.g. natural gas). When hosting your own inference on local devices this is easy to control, but for cloud-based resources it’s often more difficult to determine. An excellent tool for determining the source of power by geographic location is [ElectricityMaps.com](#). When choosing an inference provider, organisations should inquire

28 MIT Technology Review, “AI’s energy usage and climate footprint,” May 20, 2025, <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech>.





as to where the data center is located physically. For Canada, the best locations for clean energy are B.C., Québec, and the Maritimes.

Data Management: Sovereignty & Ethical Considerations

One of the key motivations for developing Canadian-based AI is to ensure data sovereignty. True sovereignty requires an AI stack that is locally governed and physically hosted on infrastructure owned by Canadian-controlled organisations (hopefully co-operatives), ensuring absolute data jurisdiction and protection against US surveillance and vendor lock-in.

Data Privacy

Hosting AI models within Canada allows organisations to maintain control over their data, ensuring compliance with Canadian data protection laws such as the *Personal Information Protection and Electronic Documents Act* ([PIPEDA](#)), which aims to regulate the use of personal information in commercial activity by private-sector organisations. This is especially relevant because under the *American Clarifying Lawful Overseas Use of Data* ([CLOUD Act](#)) (2018), US law enforcement can compel technology companies to provide data stored on their servers regardless of whether that data is located in the U.S. or on foreign soil. Also, [section 702](#) of the *Foreign Intelligence Surveillance Act* (FISA), which was added in 2008, permits targeting non-U.S. persons located outside the United States, posing a significant risk for Canadian organisations handling sensitive data like asylum applications or health records.

Ethical Data

Most large AI models are trained on data that includes copyrighted or ethically questionable content. While training a large language model is outside the average Canadian organisation's budget, opportunities for collaboration and co-operation between public and private entities may offer solutions. Additionally there are available models that are not trained with scraped or opaque training data. One of the most prominent is [Olmo](#) from the Allen Institute (Ai2). Olmo is one of the most transparent models available. Not only does it provide an excellent end-use open-weights model, the release of the training data and training checkpoints allow organisations to fork and perform their own training at any point. Another prominent open model is [Starcoder](#). Released in early 2023, this coding model is still useful three years later for code autocompletion, and stands apart for its ethical approach to training data. Most recently, the Swiss AI Initiative released [Apertus](#), a public large language model. Launched in September 2025, Apertus is under the free and open-source Apache 2.0 license and makes everything - architecture, training data, model weights - open and transparent.

Technical Futures

The choice of hardware and software plays a critical role in the technical feasibility of a Canadian sovereign AI.

Hardware: While large tech companies dominate the AI hardware market, co-operatives can explore collaborative partnerships with hardware providers to secure better pricing and more transparent supply chains while waiting for Canadian companies to catch up.²⁹

Software: Open-weights AI models such as Llama, Mistral, and Qwen offer a viable alternative to frontier models for inference. These models can be customized, deployed, and maintained

29 ISED Canada, "What We Heard Report: Consultations on AI Compute," n.d., <https://ised-isde.canada.ca/site/ised/en/what-we-heard-report-consultations-ai-compute>.





within a Canadian context, reducing reliance on foreign platforms. For example, a co-operative might choose to use the Llama-3 model, which is available under a permissive license, allowing for greater flexibility in deployment and customization.

As AI continues to evolve, the technical feasibility of hosting a Canadian sovereign AI will depend on several factors:

Scalability: The ability to scale AI systems as demand grows while maintaining performance and efficiency. For example, a co-operative might start with a small, local AI model for internal use and later scale it to support more complex tasks, such as customer service or data analysis.

Interoperability: Ensuring that AI systems can integrate with existing tools and platforms used by co-operatives and social enterprises. For instance, an AI system might need to work with a co-operative's existing CRM or data management system, requiring careful planning and integration.

Community and Collaboration: Building a collaborative ecosystem that supports open-source development, knowledge sharing, and ethical AI practices. For example, co-operatives might collaborate with other organisations to develop shared AI tools, reduce costs, and promote ethical AI development.





Part 5: Risk Assessment

In gaining a better understanding of the feasibility of co-operative AI, we must also account for and understand the technical and social risks. Our goal is resiliency and sustainable stewardship: to de-risk our joint endeavors. This section identifies and suggests mitigation for the challenges that face the development of sovereign, values-based AI tools and services.

Technical & Regulatory Risks

Data Leakage & the Shadow AI Security Perimeter

Our research shows that over half of employees at organisations without an AI policy are already using freely available Big AI tools for work. Although it might seem benign, from a security perspective, this constitutes an unauthorized data leakage. Private organisation data, documents, and internal communications are leaving ‘secured premises’ and entering third-party systems where control is lost over its storage and use. There is a risk of private co-op data being used to train public models, particularly in the case of ‘white label’ models (e.g. products branded for a specific use, but often using corporate vendor models behind the scenes). For organisations using their own versions of ChatGPT, they’ll need to be careful to turn off any training opt-ins now and in the future. Otherwise, this could lead to sensitive co-operative information being misused, perhaps even used for future training data for private company models.

Transparent and strict data governance policies and letting workers know how to use AI tools are crucial first steps in mitigating data misuse. Additional measures include using federated learning (only sharing training data weights instead of the data itself) or on-premise training (locally stored data), and ensuring that all third-party models are audited for data usage practices.

Platform Risk & Vendor Lock-in

Relying on external corporate AI integrations is not without significant risk. If a co-op builds its own AI-driven application as a ‘wrapper’ or integration layer over a commercial AI platform (e.g., OpenAI, Google, or Anthropic), it is vulnerable to changes in the platform’s API pricing, terms of service, or availability. Using commercial AI models may lead to vendor lock-in, where the co-op becomes dependent on a single provider’s infrastructure, tools, and models. Diversifying AI dependencies, using open-source and open-weights alternatives where possible, and building in flexibility to switch platforms if needed, mitigates these risks.

While these pitfalls are evident in choosing AI tools, they are a similar concern for procuring any digital technology. For organisations in the process of selecting new software, it’s critical that their procurement guidelines include open-source products, and accept diverse suppliers. Social procurement policies are integral to creating openings for smaller organisations.

Evolving Regulatory Compliance

Governments around the world are moving quickly to regulate AI use and ensure safety for their citizens and data. This landscape is changing rapidly, and like all businesses, Canadian co-operatives need to follow developments to ensure they are compliant for their business in Canada and beyond.





How important is privacy and security when choosing an AI tool?



We asked on a scale of 1 to 5, how important is privacy and security when choosing an AI tool?

The Pan-Canadian Artificial Intelligence Strategy

Launched in 2017, this federal policy and investment framework aims to advance AI research, talent acquisition and commercialisation designed to strengthen Canada's leadership in AI. Its key pillars include:

- Commercialization: Supporting businesses to commercialize AI innovations within Canada.
- Standards: Developing standards and ethical AI practices.
- Talent and Research: Investing in academic research and training to attract/retain AI experts.

Interested co-operatives would need to partner with existing research or commercialization hubs in order to access funds available under the 'Global Innovation Clusters.' This high level strategy is due to be updated in Spring 2026, following a [30-day public consultation](#) in late 2025.

Canadian Sovereign AI Compute Strategy

Announced in 2024, the [Canadian Sovereign AI Compute Strategy](#) is a \$2 billion initiative focused on developing domestic, secure infrastructure to ensure Canadian organisations (specifically researchers and AI companies) have access to computing power without relying on foreign providers. It directly funds data centers and a complementary [AI Compute Access Fund](#) to boost local AI innovation and protect Canadian data sovereignty.

While Canada is working on several federal strategies for expedited AI development, adoption, commercialization, its regulatory landscape still lags in any enforceable laws related to the ethical development and use of AI. The controversial Artificial Intelligence and Data Act (AIDA) died during the prorogation of Parliament when Justin Trudeau stepped down as prime minister.³⁰ Compliance with AIDA would have required co-ops to ensure their AI systems are transparent, fair, and accountable. No replacement legislation has yet been proposed.

³⁰ Montreal AI Ethics Institute, "The Death of Canada's Artificial Intelligence and Data Act: What Happened and What's Next for AI Regulation in Canada," n.d., <https://montrealethics.ai/the-death-of-canadas-artificial-intelligence-and-data-act-what-happened-and-whats-next-for-ai-regulation-in-canada/>.





Europe: The EU AI Act

The [EU AI Act](#) is a comprehensive regulatory framework that classifies AI systems based on their level of risk and imposes strict requirements on high-risk systems. While the Act applies to organisations operating within the EU, it may also have implications for Canadian co-ops that export services or data to the EU. Co-operatives can sign up for information on the evolving technical standards and guidelines on the [AI Act Single Information Platform](#). The Canadian Trade Commissioner's office also has [relevant information](#).

Whether your work is based in Canada or abroad, co-ops need to stay informed on regulatory developments and engage legal and compliance teams early in the AI adoption process. In particular, be sure risk assessments are done to determine if an AI system falls under the EU AI Act's scope.

Ethical & Reputational Risks

A co-operative AI product will only be as strong as its alignment with co-operative values. If we fail here, we risk 'ethics washing' and simply rebranding the same extractive infrastructure that the sector aims to counter.

Ethics Washing & Integrity

The AI tools, specifically LLMs, that currently dominate the commercial market are built on the uncompensated extraction of human labour and the mass ingestion of copyrighted material. For example, if the AI system is trained using data tagged by underpaid workers in the global South, it could undermine a co-op's values of fairness, equity, and social responsibility. A concrete example of this practice, first reported by Time magazine, was the case with Kenyan workers earning less than \$2 an hour to filter toxic content for ChatGPT.³¹ Building 'sovereign' tools on top of these foundations could undermine our reputation as a values-led sector. While there is a movement to build LLMs trained exclusively on public domain or legally licensed works, these alternatives are currently few in number and lack the capability seen in the commercial AI market. At a minimum, conducting supply chain audits to understand how data are sourced would help mitigate exploitation of workers, and help promote more transparent development practices.

Bias & Liability

AI systems inherit biases from the data they are trained on, which can lead to unfair or harmful outcomes. This issue has been widely studied (see the work of the [Distributed AI Research Institute](#) (DAIR), the [Algorithmic Justice League](#) and others), and is part of a larger concern about data colonialism and technical inequity. Cooperatively developed AI tools must account for this challenge by seeking out fully open-source models, and/or by integrating open-source auditing frameworks directly into the Solidarity Stack. There are currently several freely available tools for this auditing; such as [IBM's AI Fairness 360](#) (AIF360) and [Giskard](#), an open-source framework allowing data scientists to scan models for various vulnerabilities, including bias and performance issues. Ideally, these tools force AI inference to operate within predefined safety boundaries before the user sees the response. Ensuring AI systems are regularly audited for fairness and transparency is integral to reducing the liability risk inherent in models' responses, and should be part of every sovereign AI application.

31 Time, "OpenAI Used Kenyan Workers Earning Less Than \$2 Per Hour to Make ChatGPT Less Toxic," Jan. 18, 2023, <https://time.com/6247678/openai-chatgpt-kenya-workers/>.





Job Displacement vs. Augmentation

There is a risk that the AI tool could replace co-op workers rather than augment their roles. This could lead to job losses, reduced morale, and a loss of institutional knowledge (human wisdom). These are real, valid fears workers have right now; [existential fears](#). We must be proactive in ensuring and reassuring workers and members that these tools have their limitations and will not replace the human judgment and empathy that are intrinsic to the co-operative sector. We want to use AI tools to lighten the load of some digital tasks, such as summarizing long meeting minutes or searching through massive file systems.

To mitigate these challenges, co-operative AI tools should be developed thoughtfully, with humans at the centre of workflows and processes. For organisations adopting these applications, workers must be involved in the AI implementation, as well as offered training and support to help them adapt to new tools.

From Digital Dependency to Co-operative Autonomy

The AI landscape is evolving at a staggering speed, and the ability for values-aligned organisations to influence their technological future is narrowing. We are witnessing an unimpeded race among a few global players to achieve Artificial General Intelligence (AGI) and eventually Artificial Super Intelligence (ASI). We need to act now to ensure we are not excluded from decisions that will affect our digital lives forever.

Why the urgency? Major AI companies have begun disbanding their internal ethics boards,³² signaling a move away from accountability in a race for market dominance. Coupled with this, the CLOUD and FISA Acts mean that Canadian organisations using commercial AI platforms are already exposed to U.S. federal surveillance legislation. This creates immediate risks and compliance issues that only a sanctioned, safe alternative can solve. In addition to these challenges, Big Tech companies are currently hoarding memory chips and buying up hardware directly from manufacturers. This is driving prices to hyperinflation levels and threatening to cut smaller co-operative buyers out of the market entirely if we do not establish our own infrastructure now.

In 2026, AI is both a technology shift and an economic necessity. Alongside this will come regulatory requirements for data management, especially pertaining to AI hosted in foreign jurisdictions. The Solidarity Stack can be seen as a kind of cooperatively developed insurance policy, offering transparency to meet compliance requirements and sovereignty over vendor lock-in.

Sovereignty Roadmap

How do we make Canadian AI sovereignty a reality? A three-phase roadmap can help co-operatives take back control of their technical future.

32 Stephen B. Klein, “There is no AI ethics in tech,” LinkedIn, n.d.; Wall Street Journal, “OpenAI executive who opposed ‘adult mode’ fired for sexual discrimination,” n.d., <https://www.wsj.com/tech/ai/openai-executive-who-opposed-adult-mode-fired-for-sexual-discrimination-3159c61b>.





Phase 1: Local First Resilience & Research

Move your inference needs off U.S. providers and host locally. Educate your staff and develop internal policies.

Organisations can begin the transition to a more sovereign stack by replacing cloud-hosted AI with local inference on consumer-grade laptops, or dedicated local servers (like a Mac Mini or a PC with a GPU) for common tasks. This means hosting robust open-source LLMs on local hardware networked within the organisation, and accessed with chat interfaces similar to ChatGPT or Claude, but lacking some of the more generalised functionality.

Combined with clear and robust internal AI policies, this private LLM hosting meets the goal of secure, private chat and internal knowledge management. Feasibly, it allows for total privacy and avoids recurring cloud costs, though organisations must accept the cost of setup, speed challenges, and occasional downtime. Options such as [AnythingLLM](#) will run on most newer personal computers, but organisations with older hardware may need to begin by researching AI options outside of the large platforms and then move directly to Phase 2.

Education and policy work is also critical in this first phase. Developing an organisation-wide framework for AI adoption and usage, reviewing data handling requirements, as well as offering dedicated training on how AI works, its pros and cons, and how it fits into the values-driven co-operative context are necessary first steps.

Phase 2: Managed AI Hosting in Canada

Enterprise needs-high utilisation and uptime-require enterprise infrastructure purchased from Canadian providers. This infrastructure requires setup, management, and comes with real costs attached.

Maintaining 100% uptime for production AI hosting is difficult on consumer hardware. Local infrastructure is vulnerable to power outages, network disruptions, IP changes, and scheduled maintenance windows, all of which translate to service interruptions. This is the classic on-premises versus cloud trade-off, and most organisations have already resolved it in favour of the cloud for website and business data hosting, precisely because the operational complexity of self-managing infrastructure is high.

The most reliable option for developing co-operative AI infrastructure is renting enterprise-grade GPU hardware in the cloud, but this comes with some drawbacks. Cloud providers offer Virtual Private Servers with GPUs. However the costs are substantial, and without efficient parallelism and request batching, which is hard to implement and not especially cost effective, this approach is out of reach for many organisations. While there are some Canadian-owned companies providing these cloud services (examples include [Telus](#) and [Micrologic](#)), they are few and far between, and confirming that they offer a 100% Canadian solution requires due diligence.

The more practical starting point is purchasing LLM inference through an API endpoint rather than managing GPU infrastructure directly. In this case, the provider handles availability, scaling, and maintenance; the organisation simply connects to the endpoint and pays per token rather than per hour. For projects without consistently high utilisation, this is almost always cheaper. Importantly, sovereign API endpoints built on open-weights models are beginning to emerge, offering organisations greater control over their data and supply chain without the infrastructure burden.





As needs mature, organisations can graduate to shared, professionally managed co-operative infrastructure, for example, incorporating mission-aware assistants built with Retrieval-Augmented Generation (RAG) and targeted fine-tuning. These options are being rolled out into the Canadian context; for example, Hypha's [Rook](#) tool, and the recently launched [Augure AI](#). This requires meaningful investment in collaborative hosting, but it delivers something the cheaper alternatives cannot: a reliable, values-aligned tool that genuinely reflects the organisation's priorities and can grow with them.

Phase 3: The Canadian Solidarity Stack

Cooperatively fund and build the maintainable Canadian-based infrastructure our organisations need.

Interestingly, AI's digital enclosure has opened the door to new conversations about supporting and participating in a robust ecosystem of permanent public goods. Ideally, we want to ensure we have a lasting public resource for Canadian's computing needs. In our view, access to clean compute can be seen as the same basic right as access to clean water and air. A Solidarity Stack, one that is open-source, locally governed, and ethically audited, is the only way to ensure that AI serves the common good rather than just the bottom line of a few proprietary, likely U.S.-owned, platforms.

By considering all of the layers that comprise AI, and understanding their origins and their negative externalities, the co-operative sector can move toward supporting technology that reflects values such as self-help, democracy, and equality. An eventual goal might be full agentic capabilities and access to sovereign foundation models not trained on stolen intellectual property; in this pursuit we would be following the lead of international initiatives like the [Swiss AI Initiative](#) or India's Sovereign AI Ecosystem ([Sarvam AI](#)) to create transparent, ethically audited, and multilingual models available to the public.

In the interim, the co-operative sector should begin building the stack we need. To start this movement, we are calling for commitments for the initial capitalisation of \$400,000 by the end of 2026. We're also looking for forward-thinking anchor clients to participate as partners in this bold vision.

The stakes could not be higher. As AI becomes an increasingly foundational layer of economic and social life, decisions made now- about who builds it, who governs it, and who benefits from it—will shape the conditions of Canadian life for generations to come. The co-operative sector is uniquely positioned to lead this conversation, because we have always understood that infrastructure and labour are never neutral. Just as co-ops once built credit unions to democratize access to capital, and producer co-operatives leverage their collective labour for shared benefit, the co-operative movement today can champion a vision of AI that is in service of people and planet, not just profit. This is not a distant or abstract ambition: it begins with the choices we make today about the tools we adopt, the platforms we fund, and the values we encode into the systems that will increasingly shape our world. By owning our infrastructure, we can take one crucial step toward distributing AI more equitably, and building a stack we can trust.





AI Disclosure Statement

This feasibility report was developed with the assistance of generative AI (GenAI) tools. The following details the extent and nature of AI involvement:

AI Tools Used: Google Gemini 3, Claude Sonnet 4.6, QWen3, Nano Banana 2.

Purpose of Use: Generating initial outlines, research, rewording, argument analysis.

Specific Contributions: The Glossary was initially generated by Gemini 3. Wording suggestions and argument analysis were supported by both Gemini and Claude.

Human Oversight: All AI outputs were reviewed, edited, and verified by the authors for accuracy, originality, and relevance. No AI-generated content was submitted verbatim without substantial modification.

Ethical Considerations: This usage aligns with our AI policies and guidelines. The authors take full responsibility for the final content and attests that it represents their own intellectual effort.

Methodology

This study draws on three sources: an online survey, a qualitative interview, and desk research. The survey was distributed by email to organisations across the Canadian solidarity economy between November 2025 and January 2026, using a snowball sampling approach in which recipients were encouraged to share the survey within their networks, and received 54 responses. One in-depth interview was conducted to supplement and contextualise the survey findings. Additional context and supporting evidence was gathered through desk research. Given the snowball sampling method, survey results are not statistically representative of the sector as a whole, but are intended to provide directional insight into attitudes and practices around AI adoption.

Glossary

- **API Debt:** The cumulative technical and operational burden created when an organization becomes overly dependent on a third-party API that may change its terms, pricing, or functionality without notice.
- **API Endpoint:** In the context of AI, this refers to a specific URL where an AI service receives requests and sends responses, allowing applications to interact with AI models or functionalities.
- **API Wrapper:** A software layer that provides a simplified user interface for a complex third-party AI service. In the context of this study, we suggest that some 'Canadian' AI services are actually just wrappers around American APIs.
- **Batching and Parallelism:** Techniques used in production environments to serve multiple users simultaneously by processing multiple queries in parallel batches (batching) and multiple batches at the same time (parallelism).
- **Data Sovereignty:** The principle that an organization or nation maintains absolute legal jurisdiction and physical control over its data, protecting it from foreign surveillance (such as the U.S. CLOUD Act) or vendor lock-in.





- **Federated Learning:** A method of training AI where the training happens across multiple decentralized devices or servers (nodes), with the raw data never leaving its original location.
- **Fine-tuning:** The process of taking a pre-trained model and further training it on a specific, smaller dataset to optimize its performance for a particular task or industry context.
- **GPU (Graphics Processing Unit):** Specialized high-performance hardware required for the heavy mathematical lifting needed to train and run AI models efficiently.
- **High Availability:** A system design approach that ensures a service (like a hosted AI) remains operational and accessible with minimal downtime.
- **Inference / Inference Engine:** The process of a trained AI model providing a response to a user query. An “inference engine” (e.g., Ollama or LMStudio) is the software that manages this process on specific hardware.
- **LLMs (Large Language Models):** Advanced AI systems trained on massive datasets to understand, generate, and manipulate human language or code.
- **Local or Local-first:** A technical architecture where data processing and storage happen primarily on the user’s device or local network rather than in a central corporate cloud, ensuring higher privacy and offline resilience.
- **MCP (Model Context Protocol):** A standardized protocol that allows organisations to swap between different underlying AI models while keeping their interface and task-specific settings consistent.
- **On-premise (On-prem):** Computing infrastructure located physically within the offices or data centres of the organisation using it, rather than being hosted by a third-party cloud provider.
- **Open-weights Models:** AI models (like Llama 3 or Mistral) where the trained ‘weights’—the internal parameters that determine how the model processes information—are publicly available for independent hosting and customization.
- **PII (Personally Identifiable Information):** Sensitive data that can be used to uniquely identify an individual (e.g., names, health records, or ID numbers), requiring high levels of security.
- **Quantization:** A technique used to compress AI models by reducing the precision of their weights, allowing them to run on hardware with less VRAM (like a standard laptop) with minimal loss in accuracy.
- **RAG (Retrieval-Augmented Generation):** A technique that connects an AI model to an external, secure knowledge base (like an organization’s internal documents). This allows the AI to provide accurate, cited responses based on specific data.
- **Shadow AI:** The unauthorized use of commercial AI tools (like ChatGPT or Claude) by employees within an organization that lacks a formal AI usage policy, often leading to accidental data leakage.
- **SLMs (Small Language Models):** Smaller, less computationally intensive versions of LLMs, designed to run on consumer-grade or local hardware while maintaining high data privacy.





- **Token Tax / Token Costs:** The recurring linear costs associated with using cloud-based AI services, where users are charged based on the number of “tokens” (fragments of words) processed.
- **VRAM (Video Random Access Memory):** The specialized high-speed memory on a GPU that determines how large of an AI model can be loaded and run at one time.

Appendix A

Access the raw data from our hands-on energy testing. Download the data here: <https://cooperativeblueprint.ai/read-the-study/#energy-monitor-graphs>

Bibliography

Business Insider. “Private LLM growth expected as enterprises shift GenAI from experiments to secure domain-specific systems.” 2026. <https://markets.businessinsider.com/news/stocks/private-llm-growth-expected-as-enterprises-shift-genai-from-experiments-to-secure-domain-specific-systems-1035769007>.

Canadian Centre for the Study of Co-operatives. “The 2024 Top Co-op Issues Survey Report.” 2024. <https://usaskstudies.coop/documents/research-reports/2024-top-co-op-issues-survey-report.pdf>.

CanTrust Hosting Co-operative and Hypha Worker Co-operative. Co-operative AI Blueprint Survey. November 2025.

Co-operatives and Mutuels Canada. “The Economic Impact of Canadian Co-operatives and Mutuels.” Nov. 20, 2024. <https://canada.coop/en/the-economic-impact-of-canadian-co-operatives-and-mutuels/>.

Co-operatives and Mutuels Canada. “What New National Research Reveals About Canadians and Stability.” March 25, 2026. <https://canada.coop/en/what-new-national-research-reveals-about-canadians-and-stability/>.

Compétence Culture. “L’IA en culture : Mieux comprendre pour agir ensemble.” Nov. 21, 2025. https://competenceculture.ca/wp-content/uploads/sites/2/2025/11/ia-etude-21-novembre-2025_competence_culture.pdf.

Garriga Miret, Mònica, David Gómez Fontanills, Xavier Martínez Serrano and Enric Senabre Hidalgo. “Social and solidarity economy and shared knowledge: Facilitating and integrating the collaborative adoption of digital commons in Catalan coops.” Internet Policy Review, Feb. 6, 2026. <https://policyreview.info/articles/analysis/digital-commons-in-Catalan-coops>.

ISED Canada. “What We Heard Report: Consultations on AI Compute.” n.d. <https://ised-isde.canada.ca/site/ised/en/what-we-heard-report-consultations-ai-compute>.

Klein, Stephen B. “There is no AI ethics in tech.” LinkedIn, n.d.

MIT Technology Review. “AI’s energy usage and climate footprint.” May 20, 2025. <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech>.





Montreal AI Ethics Institute. "The Death of Canada's Artificial Intelligence and Data Act: What Happened and What's Next for AI Regulation in Canada." n.d. <https://montrealetics.ai/the-death-of-canadas-artificial-intelligence-and-data-act-what-happened-and-whats-next-for-ai-regulation-in-canada/>.

Mozilla Foundation. "A dispatch from the India AI Summit floor." Mozilla Foundation Newsletter, March 4, 2026.

One Hundred Nights. "Institutional memory and AI." n.d. <https://onehundrednights.com/article/institutional-memory-ai/>.

Platform Cooperativism Consortium. "The Solidarity Stack." Dec. 21, 2025. <https://platform.coop/blog/the-solidarity-stack/>.

Playback. "ACTRA warns CRTC of AI threat to on-screen performers." May 23, 2025. <https://playbackonline.ca/2025/05/23/actra-warns-crtc-of-ai-threat-to-on-screen-performers/>.

Population Matters. "Beyond profit: An introduction to the solidarity economy." November 2024. <https://populationmatters.org/news/2024/11/beyond-profit-an-introduction-to-the-solidarity-economy/>.

Rowinski, Michael. "Shadow AI: the hidden threat quietly undermining your business." Mimecast, April 21, 2026. <https://www.mimecast.com/blog/shadow-ai-the-hidden-threat/>.

Scholz, Trebor. "The Solidarity Stack" (keynote, Cooperative AI Conference, Istanbul, Nov. 11, 2025). Platform Cooperativism Consortium, Dec. 21, 2025. <https://platform.coop/blog/the-solidarity-stack/>. Note: Morshed Mannan is co-credited with coining this term.

SOCAN. "SOCAN's AI 'No Exceptions' campaign receives overwhelming support." n.d. <https://www.socanmagazine.ca/news/socan-ai-no-exceptions-campaign-receives-overwhelming-support/>.

Statistics Canada. "Canadian Non-Financial Co-operatives, 2021." The Daily, Nov. 6, 2023. <https://www150.statcan.gc.ca/n1/daily-quotidien/231106/dq231106c-eng.htm>.

Statistics Canada. "The economic impact of Canadian co-operatives and mutuals, 2021." 2021. <https://www150.statcan.gc.ca/n1/en/pub/11-627-m/11-627-m2026014-eng.pdf>.

Statistics Canada. "Workplace adoption of AI in Canada." June 2025. <https://www150.statcan.gc.ca/n1/pub/11-621-m/11-621-m2025008-eng.htm>.

Time. "OpenAI Used Kenyan Workers Earning Less Than \$2 Per Hour to Make ChatGPT Less Toxic." Jan. 18, 2023. <https://time.com/6247678/openai-chatgpt-kenya-workers/>.

Vokey, Sheila. "Reflections on a pivotal year for Central 1." Central 1, Jan. 13, 2026. https://www.central1.com/in_the_news/reflections-on-a-pivotal-year/.

Wall Street Journal. "OpenAI executive who opposed 'adult mode' fired for sexual discrimination." n.d. <https://www.wsj.com/tech/ai/openai-executive-who-opposed-adult-mode-fired-for-sexual-discrimination-3159c61b>.

