

Responsible AI Policy Framework

for the B.C. Health Sector



Ministry of
Health

Last Updated: June 2026

Table of Contents

Introduction to the Responsible AI Policy Framework	3
Territorial Acknowledgement	3
Background	3
Context	3
Scope and Applicability	4
Section One: Definitions	5
1.1 Artificial Intelligence Definitions	5
Section Two: Responsible AI Expectations and Requirements	7
2.1 AI System Lifecycle	7
2.2 Responsible AI Principles	10
2.3 AI Risk Management Framework	13
2.4 Informed Consent	13
2.5 Accountability and Responsibilities	14
2.6 Incident Management	15
Section Three: Procurement Considerations	16
1.1 Health Canada or OCIO Approved Systems	16
1.2 Third-Party AI Systems	16
1.3 Open-Source Systems	16
Conclusion	18
References	19

Introduction to the Responsible AI Policy Framework

Territorial Acknowledgement

The Province of B.C.'s Ministry of Health (Ministry) headquarters building is located on the territories of the lək̓ʷəŋən (Lekwungen) Peoples of the Songhees and Xwsepsum (Esquimalt) Nations, and we acknowledge, with gratitude, that much of our work on this document has taken place there. We also acknowledge, with respect, the inherent rights and title of the First Nations whose ancestral territories cover every inch of the province now colonially known as British Columbia, including their unextinguished rights to land, self-determination, health, and wellness.

We commit to continuing our collective journey towards Truth and Reconciliation, in pursuit of eliminating Indigenous-specific racism. We will uphold and honour Indigenous Peoples' inherent rights in all aspects of our policies and practices.

Background

The Ministry of Health recognizes the value that Artificial Intelligence (AI) can have in enhancing health-care delivery through improved patient experience and safety, improved clinical outcomes and provider experience, and decreased administrative pressures.

Action 29 of B.C.'s [Health Human Resources Strategy](#)¹ calls for the development of a strategy to guide the implementation of AI and workforce-enabling technologies. An important first step to building an effective AI strategy is ensuring that appropriate guardrails and enablers are in place. As such, the Ministry has developed this Responsible AI Policy Framework to guide the consistent adoption and implementation of AI across the health-care sector.

The Policy Framework provides standards and guidelines for the ethical and effective use of AI by the Ministry and B.C. health authorities.

Context

The Responsible AI Policy Framework has been developed in a broader context of federal, provincial, and territorial guidance, intended to respond to this rapidly evolving technology.

The Government of Canada released its [Pan-Canadian Artificial Intelligence Strategy](#)² in 2017, followed by a [guide on the use of generative AI](#)³ for use within federal institutions in 2023. While this guide is not a mandatory directive to provincial government entities, it outlines many best practices and considerations that are relevant for managing AI risks.

Provincially, the B.C. and Yukon Ombudspersons and Information and Privacy Commissioners released a joint report in 2021 entitled [Getting Ahead of the Curve](#),⁴ with recommendations around fairness and privacy for AI use in public service delivery. In 2024, the B.C. Office of the Chief Information Officer released a set of [draft artificial intelligence responsible use principles](#)⁵ for B.C. Public Service employees, followed by its [policy on the use of generative AI](#)⁶ in 2025. Similar guidance is emerging across other provinces and

territories, with a growing body of AI policies, principles, and tools aimed at promoting responsible use in public-sector contexts. In addition, B.C. is continuing to work with Doctors of BC to ensure alignment with any guidance that they publish.

Scope and Applicability

The Responsible AI Policy Framework applies to and is intended to guide the development and implementation of AI for the Ministry and B.C. health authorities.^a It establishes a set of standards and guidelines for personnel to adhere to in procuring, developing, operating, and using AI technologies responsibly within their organization and in line with existing provincial guidance and expectations.

The Policy Framework governs all aspects of AI systems, including individual components and modifications. It applies to AI systems developed by or for the Ministry and health authorities, including those procured from external vendors and those used by vendors to deliver services. It also applies to any AI systems used by Ministry staff that are not directly procured, approved, and governed by the B.C. Office of the Chief Information Officer.

Health authority policies and processes are expected to align with the overall provincial standards and guidelines set out in this document. The Policy Framework does not apply to care providers solely working in private practice. Care providers who work in both private practice and health authorities must adhere to the framework when working on behalf of a health authority.

A phased approach to compliance will be established for the Ministry and health authorities, with an initial compliance expectation of August 31, 2027. The Ministry will engage with health authorities to inform and develop the compliance approach.

Due to the evolving nature of AI and its usage, this Policy Framework will be updated regularly to reflect new developments and priorities, with explicit review milestones to maintain alignment with broader provincial and federal direction.

^a Includes regional health authorities, the Provincial Health Services Authority (PHSA), and BC Shared Health Services (BCSHS).

Section One: Definitions

1.1 Artificial Intelligence Definitions

Key definitions relevant to the Responsible AI Policy Framework include:

Artificial Intelligence	Information technology that performs tasks that would ordinarily require human intellectual capacity to accomplish, such as making sense of spoken language, learning behaviours, or solving problems. ⁷
AI System	A technological system, including all components, that processes data, particularly related to human activities, using models such as neural networks, machine learning, genetic algorithms, agents, or other techniques to generate outputs, including predictions, recommendations, decisions, or content. It performs tasks traditionally requiring human intelligence, such as learning, reasoning, problem-solving, data analysis, and decision-making, with varying levels of autonomy.
AI Model	A machine-based system that makes inferences from inputs to generate predictions, recommendations, or decisions. It is designed to operate with varying levels of autonomy and may adapt over time after deployment.
Agentic AI	A type of AI model or system capable of autonomous reasoning, planning, and execution of multi-step tasks to achieve high-level objectives.
Automated Decision Systems	A system that makes or supports decisions influencing access to services, care, employment, or other consequential outcomes.
Conversational AI	A type of AI that aims to provide more human experience to users through chatbots or voice bots that can not only understand human speech and language but can also produce natural responses.
Embedded AI	Artificial intelligence capabilities that are integrated into, enabled within, or augment a broader application, system, platform, device, or service whose primary purpose is not AI.

Generative AI	A type of AI model that generates new content such as text, audio, code, videos, or images, based on user-generated prompts, ambient signals, or other contextual information.
Machine Learning	The ability of computers to identify patterns, learn from data, and make inferences or decisions, without having been explicitly programmed to do so. ⁸
Traditional AI	AI that scores, predicts, classifies, or recognizes patterns in data, without producing new content.

Section Two: Responsible AI Expectations and Requirements

This section outlines the lifecycle of AI system development and the key governance elements that support its responsible use. It covers the AI system lifecycle; responsible AI principles; risk management; requirements for informed consent; accountability and defined responsibilities; and incident management processes.

2.1 AI System Lifecycle

The AI development lifecycle is a series of stages, from ideation to deployment and monitoring. The AI implementation lifecycle is iterative because AI solutions or products often go through continuous recalibration and evaluation to incorporate updated data and improved modelling techniques, even after implementation. Continuous reviews mitigate the risk of AI model performance degradation; enable improvements based on new clinical evidence, clinician and patient feedback, and technical breakthroughs; and ensure consistent accuracy, efficiency, and alignment with organizational needs.

These stages are not rigid or linear and may be bypassed or returned to. **Figure 1** serves as a simplified representation of the overall lifecycle and its key processes.




Figure 1. AI lifecycle for the implementation of AI systems

Each stage of the AI development lifecycle can be completed internally or procured from a vendor. In the case of procured AI systems, stages three (data processing) through seven (model testing, evaluation, validation and verification) are typically performed either directly by, or along with, the vendor. All stages of the lifecycle benefit from clinical and user-centered design, with engagement of clinicians, operational leaders, intended end-users, patients, and other interest-holders.

The Ministry and health authorities are not required to carry out each activity directly. Instead, the focus should be on conducting thorough risk assessments and managing deployed AI, as the Ministry and health authorities are still accountable for oversight. The specific objectives and actions typically executed within each stage of the AI lifecycle are described in **Figure 2**.




Figure 2. View of the AI lifecycle with the objectives and main activities executed at each stage.

2.2 Responsible AI Principles

AI development and implementation must be guided by principles that promote ethical, safe, effective, and equitable use.

The following Responsible AI Principles were developed by the Ministry in consultation with B.C.'s health authorities, the First Nations Health Authority (FNHA), and Doctors of BC (DoBC), and they align with other provincial and federal principles, regulations, and legislation. By adhering to these principles, the Ministry and the health authorities can leverage AI to improve health outcomes and service delivery; support AI literacy; and safeguard the health, well-being, and personal information of patients, providers, and staff.

Description	Commitment
Public Benefit	
Public Benefit focuses on ensuring that AI systems are developed and implemented to deliver clear benefits and/or insights for the public.	The use of AI systems must align with provincial priorities and goals that consider institutional and public benefits, prioritizing the needs of individuals and communities, with specific consideration for impacts to First Nations, Inuit, and Métis Peoples. The use of AI systems should align with the broader goals and objectives of the Ministry Service Plan to effectively support and enhance the overall B.C. health system, and positively contribute to the delivery of health care in B.C.
Indigenous Rights and Reconciliation	
Indigenous Rights and Reconciliation recognizes the inherent rights and title of B.C. First Nations and the inherent rights of all First Nations, Métis, and Inuit living in B.C., and strives to uphold Truth and Reconciliation in pursuit of eliminating Indigenous-specific racism in B.C.	AI systems must respect and uphold the Rights of Indigenous Peoples, aligning with the Declaration on the Rights of Indigenous Peoples Act.⁹ This includes recognizing the ongoing impact of colonialism and Indigenous-specific racism embedded in data, tools, and systems. Cultural safety must be prioritized, and AI initiatives involving Indigenous communities should include consultation with those communities.¹⁰ This principle supports Indigenous data sovereignty and commits to working towards the First Nations Principles of OCAP®¹¹ and the Métis Principles of OCAS.¹²

Accountability

Accountability refers to clear ownership and responsibility throughout the entire AI management lifecycle, ensuring compliance with applicable regulations and legislation.

All partners involved in AI system development and deployment must have clearly defined roles and responsibilities.

Organizational governance structures must ensure that AI systems adhere to legal, ethical, and industry standards, incorporating clear processes to address failures or concerns. This includes strong data governance practices, which are critical for maintaining oversight and supporting compliance across the AI system lifecycle.

Safety and Human Centricity

Safety and Human Centricity prioritizes human well-being, autonomy, and dignity in the design and use of AI systems, while actively preventing physical, psychological, and social harm.

AI technologies must be designed to protect human life and promote safety at every stage of interaction. Retaining human oversight over decisions maintains trust, safeguards quality, and ensures that AI systems support, rather than replace, clinical expertise.

If patient safety is harmed and/or threatened in ways that cannot be justified or mitigated, the use of AI must immediately be stopped and reconsidered.

Equity and Anti-Discrimination

Equity and Anti-Discrimination aims to ensure that all patients and health care staff are treated equitably and without bias or discrimination.

Ministry and health authority personnel must be aware of their own perspectives, socially dominant concepts, normative ideals, and the biases that come with them.

AI systems must be designed and used responsibly to prevent the amplification of biases, systemic inequities, and discrimination; to make use of representative datasets whenever possible; to better meet the evolving realities of marginalized communities; and to accommodate non-dominant cultural and social frameworks.

AI systems used by the Ministry and the health authorities must comply with the requirements of the [Anti-Racism Data Act](#).¹³ Additionally, ongoing system testing is required to avoid the perpetuation of existing biases—such as those based on race, gender, geography, or socioeconomic status—and ensure that mechanisms exist to address and mitigate them.

Privacy and Data Security

Privacy and Data Security refers to the commitment to protect personal information, data, systems, and operations against unauthorized access, misuse, and harm, while addressing confidentiality and autonomy concerns.

Security measures must be implemented, including clear consent mechanisms, transparent communication, and proper handling of both personally identifiable and de-identified information.

The integrity of AI models must be safeguarded from external threats, in line with the requirements of current and evolving regulations. This includes ensuring data is stored securely and in accordance with provincial and national legislation, and with regard to digital and data sovereignty considerations.

Transparency

Transparency is crucial for building trust in AI systems, and refers to being open about the purpose, design, effects, and usage guidelines of AI tools.

Transparency ensures that the rationale behind AI decisions is accessible, understandable, and open to scrutiny. This includes the concepts of explainability^b and interpretability,^c meaning that AI models should be designed and communicated in ways that allow users - especially subject matter experts - to understand how outputs are generated, validate results, and challenge recommendations when needed.

Transparency supports informed decision-making by Ministry and health authorities' personnel, ultimately fostering trust and enhancing the effectiveness of AI systems.

Sustainability

AI systems in healthcare must demonstrate sustainability through fiscal responsibility, sustained performance, and reduced environmental impacts.

AI technologies must be developed and deployed in ways that promote financial sustainability, environmental responsibility, and long-term system efficiency.

AI systems must be designed to minimize environmental impacts and resource consumption, while ensuring they are resilient, scalable, and maintainable over time.

Consistent performance must be supported through reasonable upkeep, and models must remain effective across diverse patient populations and clinical scenarios.

^b Explainability is the degree to which an external observer can understand why an AI system made a specific decision or prediction.

^c Interpretability is the extent to which a human can comprehend the cause of a decision made by an AI system, often involving the ability to follow the model's reasoning process.

2.3 AI Risk Management Framework

AI use cases should be assessed during ideation and discovery to confirm that AI is the right solution for a given problem and to explore possible unintended impacts in other areas. Organizations should also periodically reassess their operational strengths and weaknesses to ensure readiness for AI adoption.

Risk assessments must be completed before and during procurement, development, implementation, and operation of any AI system. The results should inform the Ministry and health authorities' decisions on resource allocation, gap identification, and risk mitigation planning. Organizations should build on existing internal risk assessment and mitigation procedures, review relevant contract terms, and incorporate emerging evidence-based risk information and practices where appropriate.

During selection, AI systems and vendors must be assessed to ensure alignment with this Policy Framework. This requirement also applies to non-AI technologies that include embedded, optional, or vendor-enabled AI capabilities ("embedded AI"), which should be evaluated with the same rigor to confirm they meet policy expectations.

In parallel, the Ministry and health authorities must comply with all relevant laws and regulations, as well as internal privacy policies governing confidential and personal data. Personnel must not provide financial, organizational, or personally identifiable information to any unauthorized party or AI system. After implementation, AI systems and embedded AI must be continuously monitored and evaluated to mitigate risks and maintain ongoing compliance. When updates or changes affect core functionality or system performance, a new risk assessment is required to evaluate risks introduced by the modification, including changes to data processing methods, system functionality, and algorithmic behaviour.

2.4 Informed Consent

Where an AI system creates or uses personal information that can be re-identified, or where a person interacts with an AI agent, informed consent is a fundamental requirement to ensure that people are made fully aware of, and that they agree to, both the use of their data and the implications of the AI system being used. People must be informed about how their data will be collected, processed, stored, and shared, as well as the measures in place to protect their privacy and security.

The Ministry and health authorities must follow existing and relevant legislative requirements regarding consent in B.C. (e.g., [Personal Information Protection Act](#),¹⁴ [Declaration on the Rights of Indigenous Peoples Act](#), etc.), including what is applicable and necessary for opt-in and opt-out options. Organizations must prioritize transparency in their approach, ensuring that the consent process is inclusive and considering the diverse needs, languages, and understanding levels of different populations.

2.5 Accountability and Responsibilities

Ministry and health authority personnel must outline the roles and responsibilities of the parties involved in the development, implementation, management, and approval of AI systems. The accountability matrix should designate a lead system owner,^d who will be responsible and accountable for ensuring the effective implementation and ongoing management of the AI system, and identify other key roles.

Accountability and ownership of an AI system's outcomes must be clearly defined to encourage transparency. Using a human-in-the-loop approach^e may be necessary to evaluate the accuracy and validity of an output from an AI system—whether fully or partially AI-generated—before relying on it for internal or patient-related work. **Where an AI system is being used to support patient diagnosis, a human-in-the-loop is a requirement.**

Where applicable, a team of experts in AI ethics, law, privacy, and technology should work together to effectively assess and mitigate AI-related risks, such as bias, hallucinations, copyright infringement, privacy compliance, and ethics breaches.

Users of AI systems are also responsible for applying AI tools appropriately, exercising professional judgment, and ensuring that all use is ethical, safe, and in alignment with this Policy Framework. Organizations are responsible for ensuring users have AI literacy and receive appropriate AI system specific training.

System owners must create and update records relating to the performance of AI systems and store such records in approved storage locations. These records will be important to inform ongoing operations, share learning, and support future decisions. Additionally, system owners must ensure that these records are kept updated to reflect any changes made to roles and responsibilities. Records related to AI systems, including outputs, must be managed and maintained appropriately as per existing and applicable organizational records retention policies and practices.

^d A system owner refers to the individual or entity designated with the authority and responsibility for the overall management and accountability of an AI system, including its development, implementation, and ongoing performance.

^e "Human-in-the-loop" means that a person is involved in the process of checking and verifying the results produced by an AI system.

2.6 Incident Management

Organizations must establish processes that allow for incident reporting of AI systems that do not function as intended. AI incidents include AI system performance outages, as well instances where the development, use, or malfunction of one or more AI systems directly or indirectly causes harm (e.g., information and security incidents or hacking of critical infrastructure).

Ministry and health authority personnel should always refer to internal incident management processes and policies, in addition to the following directions:



- Personnel must report to management and the accountable system owner any incidents where AI systems are not performing to expectations.
- Personnel must report instances where AI usage may contradict Standards of Conduct, this Policy Framework, or applicable laws, regulations, or standards.
- Clear and accessible processes should be established and open for external users, third parties, and anyone who interacts with an AI system, to report incidents.

Section Three: Procurement Considerations

AI systems may be developed internally, procured from external vendors, or enabled through embedded or open-source components. This section outlines the additional considerations that apply at the procurement and acquisition stage, including vendor assessments, approvals, and contractual clarity. The intent is to ensure that AI systems introduced through procurement are fit for purpose, responsibly governed, and aligned with provincial expectations.

1.1 Health Canada or OCIO Approved Systems

For AI systems procured or approved by the B.C. Office of the Chief Information Officer (OCIO), or approved by Health Canada as a medical device, OCIO and Health Canada approvals address regulatory and technical compliance; however, the Ministry and health authorities are responsible for assessing operational, ethical, and patient safety risks specific to their environments.

1.2 Third-Party AI Systems

Where third-party AI systems are procured from external providers or vendors, a thorough vendor and system assessment is needed before pursuing a vendor relationship.

Third-party AI tools should only be used for approved, intended uses, which must be clarified and defined in purchase contracts, agreements, and terms of service. Use of an AI system intended for internal purposes should not be extended to patient-related uses or other parties unless explicitly stated in the licensing agreement. If an AI system is designed to support patient use, the terms of the license must clearly outline the scope of use and any applicable restrictions.

1.3 Open-Source Systems

Open-source^f AI systems are models and tools that are freely available for anyone to use, modify, and distribute. Open-source tools encourage collaboration and information sharing; however, they come with unique risks, such as limited or lack of contractual recourse. For this reason, open-source AI systems should be clearly identified throughout the risk assessment process to ensure that these unique risks are both managed and mitigated.

When new features are developed and integrated into an open-source AI system, they must be evaluated through the same risk assessment process as closed-source features. Any specific risk mitigations related to open-source systems must be diligently implemented.

^f Open-source AI refers to publicly available AI models, open-access datasets, and frameworks that allow users to access and modify the underlying code. It does not include proprietary AI systems like ChatGPT or Google Gemini, which are owned by private companies and have restrictions on their use and modification.

The Ministry and the health authorities must individually determine the contexts in which open-source AI systems may be used, at the discretion of executive leadership.

Conclusion

This Responsible AI Policy Framework sets clear standards and guidelines for the Ministry of Health and B.C. health authorities to develop, procure, operate and use AI for real public benefit.

The use of AI can support better care and stronger services for people in British Columbia, and it must be done in a way that protects people's safety, rights, and trust.

This framework will be updated over time as AI evolves to reflect new developments and priorities, and to maintain alignment with broader provincial and federal direction.

If you have questions about this Responsible AI Policy Framework, please reach out to HealthAI.Policy@gov.bc.ca.

References

- ¹ Government of British Columbia. B.C.'s health human resources strategy: putting people first. Victoria (BC): Government of British Columbia; 2022 Sep 29. Available from: <https://news.gov.bc.ca/files/BCHealthHumanResourcesStrategy-Sept2022.pdf>
- ² Government of Canada. Pan-Canadian artificial intelligence strategy [Internet]. Ottawa: Innovation, Science and Economic Development Canada; modified 2024 Dec 3 [cited 2025 May 1]. Available from: <https://ised-isde.canada.ca/site/ai-strategy/en>
- ³ Government of Canada. Guide on the use of generative artificial intelligence [Internet]. Ottawa: Government of Canada; modified 2024 Oct 15 [cited 2025 May 1]. Available from: <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/guide-use-generative-ai.html>
- ⁴ Offices of the Ombudspersons and Information and Privacy Commissioners of British Columbia and Yukon. Getting ahead of the curve: meeting the challenges to privacy and fairness arising from the use of artificial intelligence in the public sector. Joint special report no. 2. Victoria (BC); Whitehorse (YT): Offices of the Ombudspersons and Information and Privacy Commissioners of British Columbia and Yukon; 2021 Jun. Available from: <https://bcombudsperson.ca/assets/media/OMB-GettingAheadoftheCurve-v6.pdf>
- ⁵ Government of British Columbia. Draft artificial intelligence responsible use principles [Internet]. Victoria (BC): Office of the Chief Information Officer; 2024 Nov 26 [cited 2025 May 1]. Available from: <https://digital.gov.bc.ca/policies-standards/ai/draft-responsible-use-principles/>
- ⁶ Government of British Columbia. Policy on the use of generative AI [Internet]. Victoria (BC): Office of the Chief Information Officer; 2025 May 9 [cited 2025 May 29]. Available from: <https://digital.gov.bc.ca/policies-standards/ai/policy/>
- ⁷ Government of Canada. Directive on automated decision-making [Internet]. Ottawa: Treasury Board of Canada Secretariat; modified 2023 Apr 25 [cited 2025 May 1]. Available from: <https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32592>
- ⁸ Government of Canada. Machine learning [Internet]. Ottawa: National Research Council Canada; modified 2023 Nov 8 [cited 2025 May 1]. Available from: <https://nrc.canada.ca/en/research-development/products-services/technical-advisory-services/machine-learning>
- ⁹ British Columbia. *Declaration on the Rights of Indigenous Peoples Act, SBC 2019, c. 44*. Victoria (BC): King's Printer; 2019 Nov 28. Available from: <https://www.bclaws.gov.bc.ca/civix/document/id/complete/statreg/19044>
- ¹⁰ Lewis JE, Abdilla A, Arista N, et al. Indigenous protocol and artificial intelligence position paper [Internet]. Indigenous Protocol and Artificial Intelligence Working Group; Canadian Institute for Advanced Research; 2020 [cited 2025 May 1]. Available from: <https://www.indigenous-ai.net/position-paper/>
- ¹¹ First Nations Information Governance Centre. The Fundamentals of OCAP® [Internet]. Akwesasne (ON): First Nations Information Governance Centre; [cited 2025 May 1]. Available from: <https://fnigc.ca/ocap-training/>
- ¹² Strayer C, Grebliunas R. Métis governance practices [Internet]. In: Indigenous digital literacies. Victoria (BC): BC campus; 2024 [cited 2025 May 1]. Available from: <https://opentextbc.ca/indigenoussdigitalliteracies/chapter/metis-governance/>
- ¹³ British Columbia. *Anti-Racism Data Act, SBC 2022, c. 18*. Victoria (BC): King's Printer; 2022 Jun 2. Available from: <https://www.bclaws.gov.bc.ca/civix/document/id/complete/statreg/22018>
- ¹⁴ British Columbia. *Personal Information Protection Act, SBC 2003, c. 63*. Victoria (BC): King's Printer; 2003 Oct 23. Available from: https://www.bclaws.gov.bc.ca/civix/document/id/complete/statreg/03063_01