Governing AI in Ontario hospitals: oversight, opportunity and risk

A governance guide for Ontario Hospital Boards of Directors

July 2026

Authors: Michael Watts, Susan Newell and Sam Ip

OSLER

A note on the origins of this guide

This guide began with an accreditation mandate. I was engaged to help an Ontario hospital board prepare for its 2026 Accreditation Canada survey, benchmarking the hospital's governance instruments — its articles, by-laws, strategic plan, vision and mission, board and committee terms of reference, and principal policies, including its enterprise risk management framework — against Accreditation Canada's Qmentum standards (issued by the Health Standards Organization) and the Healthcare Insurance Reciprocal of Canada's (HIROC) standards. Neither identified artificial intelligence (AI) as a matter the board was expected to oversee or monitor. That silence prompted this guide.

I relied heavily on AI to research the field, identify best practices, locate authority, and draft. The work demonstrated both its power and AI's characteristic risk: fluent, confident output that can be wrong, including invented facts and citations to authorities that do not exist — a failure mode known as "hallucination." Every factual statement and source in this guide was verified against primary authority by its human authors. That discipline is the guide's thesis: AI can strengthen the work of a board and its advisors, but only within a framework of human verification and accountability of the kind this guide builds.

My co-authors strengthened the guide throughout: Susan Newell contributed across its development, with particular thanks for her work on the guide's analysis of Health Canada's regulatory framework for machine learning-enabled medical devices; and Sam Ip for AI congruency and the contractual framework and compliance. I am grateful to Jennifer Bieman for her writing assistance, and to the clients who reviewed earlier drafts and sharpened the guide with their practical judgment. Osler's editorial group provided a final review of the writing.

— Michael Watts

Contributors:

Michael Watts

Partner, Corporate Commercial and Chair, Health Industry Group, Toronto

Susan Newell

Partner, Health, Toronto

Sam Ip

Partner, Technology, Toronto

Table of Contents

Executive summary

Artificial intelligence is already at work in Ontario hospitals — drafting clinical documentation, flagging deteriorating patients, reading medical images, and scheduling operating rooms. Every hospital board is therefore already exercising AI oversight, deliberately or by default. No AI-specific statute is needed to engage the duty. Under the *Public Hospitals Act*, its Hospital Management Regulation and the *Not-for-Profit Corporations Act, 2010*, boards must govern and monitor the activities of the hospital, and AI is now one of those activities. The regulatory environment is moving: the first regulation under the *Enhancing Digital Security and Trust Act, 2024* — the Cyber Security Regulation, O. Reg. 51/26 — takes effect July 1, 2026, with AI-specific regulations expected to follow. Yet the signals boards traditionally rely on to surface governance priorities are silent. Neither Accreditation Canada's Qmentum standards nor HIROC's risk taxonomy names AI as a governance matter. The duty is current, and the business judgment rule protects a documented, informed process — not inaction.

This guide provides a governance architecture proportionate to hospital resources. AI is governed in two deployment streams — operational efficiency AI and clinical AI — and every deployment is classified into one of four risk tiers, with approval authority escalating with risk: CIO/CTO approval for low-risk operational tools; senior medical officer endorsement and CEO approval for clinical decision support and autonomous diagnostic AI; and a prior board resolution for agentic AI. An AI Management Committee reporting to the CEO executes deployment governance; the Audit/Finance Committee oversees the operational stream and the Quality Committee the clinical stream; the Medical Advisory Committee's statutory role is preserved. Structured reporting — a semi-annual executive report, an annual performance and compliance report, and 48-hour notification of material incidents — keeps the board informed as its monitoring duty requires. A sample board policy (Schedule D), an adoption resolution (Schedule C), report templates and contractual guidance operationalize the model.

The framework asks the board to make a small number of risk-appetite decisions: whether to adopt the policy and the four-tier framework; where to set the materiality thresholds that trigger board notification; the board's appetite for the highest-risk deployments — agentic AI requires prior board approval, and certain uses are prohibited under the policy outright; the investment it will commit to director, executive and staff AI fluency; and its vendor posture, including independent validation of high-risk clinical AI rather than reliance on vendor self-attestation.

After reading this summary, directors should review the legal foundation in Articles 3 through 5, ask management to present the current AI inventory against the four-tier framework in section 7.3, and direct completion of the 20-question board checklist at Schedule E. Boards ready to act can adopt the Schedule C resolution and the Schedule D policy, adapted to the hospital's circumstances on the advice of counsel. A board that takes these steps has begun to discharge a duty that was activated the day its hospital's first AI tool went live.

Article 1

Introduction

1.1 Overview: how to use this guide

Governing AI in Ontario Hospitals: Oversight, Opportunity, and Risk (the guide) has been prepared in order to assist Ontario public hospital (hospital) boards of directors (boards) in discharging their artificial intelligence (AI) governance oversight obligations by addressing both the risks and strategic opportunities AI presents.[1] AI in healthcare is not a static technology, but a rapidly evolving capability. Diagnostic tools learn from new data, algorithms drift as clinical practice changes, and AI systems are emerging that can autonomously execute multi-step clinical tasks without human initiation at each step.[2] A structured governance framework is essential because it provides the institutional architecture to evaluate, deploy, monitor, and adapt to AI innovations as they emerge, rather than forcing hospitals to respond to each development on an ad hoc basis. Directors pressed for time should focus on Article 1, Article 3, the conclusion and schedule E.

This guide addresses five interconnected themes.

- First, it addresses the fact that there is a legal foundation that already requires boards to govern AI as an institutional activity, even in the absence of AI-specific regulations.
- Second, the guide canvasses four risk tiers of AI deployments and the governance and operational considerations for each risk tier. Approval authority for AI deployments should be proportionate to risk, and may range from chief information officer or chief technology officer approval requirements for low-risk operational tools, to full board decisions for agentic AI.
- Third, the guide outlines a proposed structured accountability for AI governance and identifies the roles and responsibilities of various governance entities within the hospital, including the board, AI Management Committee (AMC), Medical Advisory Committee (MAC) and Quality Committee.
- Fourth, the guide highlights cybersecurity considerations and outlines existing legal obligations related to AI-enabled threats that boards must address.
- Finally, the guide analyzes the opportunities that operationally sound AI deployments provide to hospitals. Subject to available resources, it could be argued that boards that govern only risk while failing to capture AI's clinical and operational benefits have failed in their strategic governance duties.

1 The guide references source material that is current to May 25, 2026. Readers should be aware that legislation, regulations, case law, regulatory authority guidance and other legal authorities referenced herein may have been amended, repealed, or otherwise modified after this date. Readers are advised to verify the current status of all referenced authorities and to seek independent legal advice before relying on the information contained in this guide for specific legal matters.

2 See Schedule B for additional examples of AI deployments in the health sector.

Given that there is both an overwhelming amount of information and resources about AI being published at a rapid cadence, and a lack of specific guidance of practical application to boards, this guide includes proposed governance architecture that hospitals could adopt to govern their use of AI. The guide also includes risk frameworks and other practical governance tools for boards to consider, including

- a sample board AI governance policy in schedule D (the Sample Board AI Governance Policy) and related documents for hospitals procuring AI solutions from vendors, including template reporting tools
- a sample board resolution for the adoption of an AI governance framework
- a board governance checklist

1.2 The existing legislative obligations

Boards do not need AI-specific regulation to trigger their AI oversight obligations; the obligation for boards to have oversight over the use of AI in their organizations already applies under existing laws applicable to hospitals. Under the *Ontario Not-for-Profit Corporations Act, 2010* (ONCA)[3] and the *Public Hospitals Act* (PHA)[4] boards are required to govern and monitor hospital activities. AI is now one of those activities. At every hospital that has deployed it, AI is no longer a peripheral technology experiment but an institutional function, meaning that a board's fiduciary duty, duty of care, and express monitoring obligations attach to AI directly. There is limited guidance from regulatory authorities on certain AI uses, such as AI scribes, but there are no detailed laws or regulations specific to AI in force. Federally, Canada has signalled a targeted, multi-statute approach — privacy and online-safety legislation rather than a comprehensive AI Act, so hospitals should track sector-specific bills rather than await a single federal AI statute.

The *Enhancing Digital Security and Trust Act* (EDSTA)[5] establishes the legislative architecture for AI accountability by enabling the regulation of the use of AI systems and cybersecurity by the Ontario public sector (OPS), Crown corporations and the broader public sector (i.e., hospitals). Enabling regulations to the EDSTA are not yet in force; however, hospitals should build the required framework now so that compliance is achieved before enforcement begins and understand that additional regulations may be issued under the EDSTA in the future.

1.3 Overview of AI streams

This guide addresses AI in two deployment streams based on the purposes for which AI is utilized: stream 1 (operational efficiency AI) and stream 2 (clinical AI). Each deployment stream has different governance imperatives. The risks associated with AI fall into two key categories: (a) risks arising from the hospital's own deployment of AI in streams 1 and 2; and (b) cyber risks from both external adversaries and the cyber exposure created by AI systems the hospital deploys.

3 *Ontario Not-for-Profit Corporations Act, 2010, SO 2010*, c 15, s. 21 (management and supervision); s. 43(1) (fiduciary duty and duty of care); s. 43(2) (duty to comply with Act and by-laws); s. 43(3) (duties non-waivable); s. 44 (reasonable diligence defence) [*ONCA*].

4 *Public Hospitals Act*, RSO 1990, c P.40 [PHA]; O. Reg. 965, ss. 2(1) and 2(3) [O. Reg. 965].

5 *Strengthening Cyber Security and Building Trust in the Public Sector Act*, 2024, SO 2024, c 26; Schedule 1: *Enhancing Digital Security and Trust Act*; Royal Assent received November 25, 2024 [*EDSTA*].

(a) Stream 1: operational efficiency AI

Stream 1 encompasses tools that optimize patient flow, scheduling, bed allocation, staffing, supply chain, and administrative workflows without direct clinical decision-making. This type of AI use by hospitals has direct implications for the fiscal imperative discussed at section 2.2. AI-optimized surgical scheduling, predictive discharge planning, and emergency department volume forecasting are mechanisms that could assist hospitals in closing the gap between what the province will fund and what patients need. These tools carry lower governance risk than stream 2 AI uses and consideration should be given to implementing them promptly.

(b) Stream 2: clinical AI

Stream 2 encompasses tools with direct clinical interfaces such as patient deterioration prediction, diagnostic imaging AI, clinical decision support, AI scribes, patient-facing AI, and research AI. These tools require evidence-based governance, phased deployment, and clinical oversight. Some of these tools may also be considered software as a medical device (SaMD) that is governed by the *Food and Drugs Act* and Medical Devices Regulations thereunder,[6] and, in some circumstances, may require Health Canada authorization before being imported or sold in Canada, as discussed in Article 6.

(c) Risk categories: AI deployment risks and cyber risks

The risks AI presents to hospitals fall into two key categories. The first is the risk arising from the hospital's own deployment of AI in streams 1 and 2, including patient safety, care quality, privacy, bias, automation bias, accountability, vendor risks and algorithm drift, the degradation of AI performance over time as the circumstances in which AI is deployed diverge from an unchanged, underlying algorithm. The second is cyber risk.

Cyber risk itself has two dimensions:

(a) AI-enabled cyber threats from external adversaries who use AI capabilities against the hospital's infrastructure regardless of whether the hospital deploys any AI of its own; and

(b) the additional cyber exposure created by the AI systems the hospital itself deploys, including expanded attack surface and new vendor and supply-chain dependencies.

Boards should consider both dimensions of cyber risk together with deployment risk and confirm that the AI systems the hospital deploys are protected from compromise and that the hospital's people and processes are positioned to use AI safely.

1.4 Methodology

This guide uses the EDSTA definition of "artificial intelligence system": a machine-based system that, for explicit or implicit objectives, infers from input it receives to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments, and such other systems as may be prescribed.[7]

6 *Food and Drugs Act*, RSC 1985, c F-27 [*Food and Drugs Act*]; Medical Devices Regulations, SOR/98-282 [*Medical Devices Regulations*].

7 EDSTA, s. 1(1).

This guide applies a proportionality principle to AI governance. While AI governance obligations are mandatory for every hospital, the depth of implementation is calibrated to institutional size, complexity and AI deployment profile. A large, urban academic health sciences centre and a community hospital in northern Ontario may build the same governance components at different scales.

For hospitals that are engaged in their own AI-enabled product development, there are additional considerations relevant to establishing appropriate oversight and governance processes, risk management and lifecycle practices that have not been addressed in this guide or the Sample Board AI Governance Policy. Although this guide does address certain AI-specific privacy considerations, it does not address all potentially relevant considerations under applicable privacy laws.

Article 2

The AI landscape: what every director must understand

2.1 AI is not coming; it has arrived

AI has moved from promise to practice in hospitals. Across the province, ambient AI systems quietly document patient-clinician encounters in real time.[8] In a large Toronto hospital, predictive tools scan for early signs of deterioration before clinicians might otherwise notice them.[9] AI-enhanced diagnostic imaging for breast and prostate cancer flags radiological abnormalities with speed and consistency that augment clinical judgment.[10] Outside the clinical setting, an AI-enhanced medical device developed by Toronto's University Health Network for remote care and patient self-management is reducing repeat hospitalizations for heart failure.[11] Behind the scenes, operational AI is reshaping how hospitals function: optimizing patient flow, surgical scheduling, bed management, and staffing in ways that were previously impossible at scale.[12] AI tools are enabling boards to manage meetings, publish minutes and ensure transparency.[13] From relatively simple tools, such as AI scribes, to complex ecosystems of AI-enabled diagnostic and clinical decision-making platforms, these developments illustrate that AI is no longer a peripheral experiment in Ontario's health care system, but an increasingly embedded feature of how care is delivered, managed, and governed.

8 See, for example: "Ontario AI Scribe Program: Helping clinicians spend more time with patients", Supply Ontario, (February 10, 2026). See also "Interim Report - Evaluating Artificial Intelligence (AI) Scribes to Reduce Administrative Burden and Enhance Data Interoperability in Canadian Primary Care", Canada Health Infoway, (November 3, 2025).

9 "Clinical evaluation of a machine learning–based early warning system for patient deterioration", Verma A et al. *Canadian Medical Association Journal* (September 16, 2024).

10 "Artificial intelligence-supported screen reading versus standard double reading in the Mammography Screening with Artificial Intelligence trial (MASAI): a clinical safety analysis of a randomised, controlled, non-inferiority, single-blinded, screening accuracy study", Lang M et al., *Lancet Oncology* (August 2023). The MASAI trial established AI-assisted mammography was non-inferior to standard double reading in mammography screening. AI-supported screening among 39,996 participants resulted in 244 screen-detected cancers. Standard screening among 40,024 participants resulted in 203 screen-detected cancers; See also "Artificial intelligence and radiologists in prostate cancer detection on MRI (PI-CAI): an international, paired, non-inferiority, confirmatory study" Saha A et al. *Lancet Oncology* (July 2024).

11 See "The Effect of Using a Remote Patient Management Platform in Optimizing Guideline-Directed Medical Therapy in Heart Failure Patients: A Randomized Controlled Trial", Science Direct. (April 1, 2024). Medly, an AI-enhanced medical device for remote care and patient self-management for individuals with heart failure, has significantly reduced heart failure hospitalizations; See also "Medly: A Digital Health Platform Taking Care Into the Future", UHN Commercialization (accessed May 2026)

12 See, for example: "Unity Health Toronto developed AI tools to ease administrative burdens. New funding is helping to scale them", Unity Health Toronto, (October 24, 2024).

13 See, for example: Diligent AI (2026).

According to Canada's National Artificial Intelligence Strategy: AI for All dated June 4, 2026, Canada faces a significant AI adoption and trust gap. Only 12 percent of Canadian businesses currently use AI, and roughly eight percent of SMEs have adopted it — well behind Nordic countries, Germany, and France. Compounding this gap, Canada ranks 44 of 47 countries on AI training and literacy and 42 on trust in AI systems. Low adoption, low literacy, and low trust are preventing Canadians from realizing AI's potential benefits.[14]

The federal government has identified the health and life sciences sectors as a priority focal point for concentrating strategic investment for Canada to build and hold a global leadership position in applying AI to improve health outcomes and deliver faster, more responsive health care services for Canadians.[15]

2.2 The fiscal imperative

For Ontario boards, AI governance is no longer a strategic choice. The deployment of AI to achieve operational efficiencies is an institutional necessity to address funding challenges. Ontario's Ministry of Health has directed hospitals to submit three-year balanced budget plans premised on 2% annual funding increases, against a Financial Accountability Office (FAO) finding that 4% is required to maintain current service levels.[16] The resulting structural funding gap — $3.4 billion in 2025-26, rising to $9.6 billion by 2027-2028 — means hospitals must deliver the same or better care with fewer resources.[17]

2.3 The governance advantage: why building first wins

Hospitals that have established robust AI governance frameworks, and related safeguards, are better positioned to realize the potential of AI-enabled operational efficiencies than institutions without a clear and consistent approach to AI governance. Hospitals without such frameworks are forced to evaluate rapidly evolving AI innovations on a case-by-case basis, without a comprehensive system to evaluate, implement and monitor the AI systems.

Among U.S. hospitals using predictive AI, 86% of multi-hospital systems reported deployment versus only 37% of independent facilities.[18] This AI adoption gap mirrors the divide between Ontario's larger academic and community hospitals and smaller rural and northern facilities. The gap in AI adoption may exacerbate existing disparities between large and small hospitals in clinical service offerings, care quality, patient safety and operational efficiencies.[19] Institutional access

14 "Canada's National Artificial Intelligence Strategy: AI for All", Government of Canada, (June 4, 2026).

15 *Ibid.*

16 "Ontario hospitals told to find savings; ministry doesn't rule out service cuts", The Canadian Press, (October 20, 2025). "Ontario Health Sector: 2025 Spending Plan Review" [*FAO Report*]. Financial Accountability Office of Ontario, (October 23, 2025).

17 FAO Report at 1.

18 "Hospital Trends in the Use, Evaluation, and Governance of Predictive AI, 2023-2024 - ONC Health IT Research & Analysis", Chang W, Owusu-Mensah P, Everson J, Richwine C,. ASTP Health IT Data Brief. Washington (DC): Office of the Assistant Secretary for Technology Policy, (September 2025).

19 The access divide already documented in Canadian healthcare, resulting in survival differences of up to 19.2% between urban and rural centres for out-of-hospital cardiac arrest. See "Urban-Rural Differences in Cardiac Arrest Outcomes: A Retrospective Population-Based Cohort Study" Connolly M, Goldstein J et al., PubMed, (December 30, 2021).

to AI technology creates a two-pathway risk that boards must understand. On the first pathway, hospitals that cannot access AI or lack the resources and talent to deploy it will struggle with the funding gap described in section 2.2. For instance, a significant number of Ontario's smaller, rural, and northern hospitals lack the financial resources, technical infrastructure, or internal expertise to deploy AI. Some lack a functioning electronic medical record (EMR) and are operating using paper records. Boards with limited resources must factor this reality into strategic resource allocation.

2.4 The OPS AI framework and AI directive

In September 2023, Ontario's Ministry of Public and Business Service Delivery and Procurement published the OPS AI Framework, which defines the OPS approach to the use of AI.[20] The framework includes the AI directive, which establishes requirements for ministries and provincial agencies that use AI systems. The AI directive is anchored in six principles intended to guide the ethical, transparent and reliable use of AI: (a) AI is used to benefit the people of Ontario (b) AI use is justified and proportionate, and AI systems used are reliable and valid (c) AI is used in a safe, secure and protective way (d) AI use is human rights-affirming and non-discriminatory (e) AI use is transparent, and meaningful explanations of decisions are made available and (f) AI use is accountable and responsible.[21]

While the OPS AI Framework and AI directive apply directly to OPS ministries and provincial agencies rather than hospitals, they reflect the responsible-use principles, governance disciplines and project-level review mechanisms that public-sector entities, including hospitals, are increasingly expected to operationalize. Boards should treat the OPS AI Framework and AI directive as informative reference points when designing their own AI governance, risk-assessment and project-approval processes.

20 "Ontario's Trustworthy Artificial Intelligence (AI) Framework", Government of Ontario. Published September 14, 2023; Updated December 1, 2025.

21 "Responsible Use of Artificial Intelligence Directive", Government of Ontario. Published January 7, 2022; Updated January 7, 2026. See also "Use of Artificial Intelligence in the Ontario Government: Independent Auditor's Report". Office of the Auditor General of Ontario, (May 12, 2026).

2.5 Key areas of regulation-making authority under EDSTA

The Office of the Auditor General of Ontario's May 2026 report, *Use of Artificial Intelligence in the Ontario Government* (Auditor General's AI Report) summarizes the five key areas regulations that are made pursuant to the EDSTA are expected to address:[22]

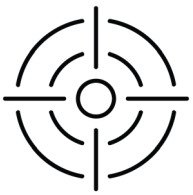	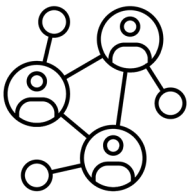			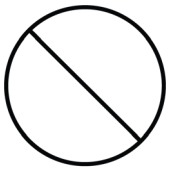
Risk Management	**Accountability framework**	**Transparency**	**Compliance**	**Prohibition**
EDSTA regulations may require public sector entities to take prescribed steps to manage risks associated with AI system use.	Regulations pursuant to the EDSTA may require public sector entities to implement AI system accountability frameworks.	Regulations may require public sector entities to provide information to the public regarding the use of AI systems.	AI systems must be used in accordance with requirements.	Prohibition. Prevent AI misuse under future regulations.

These five areas frame the substantive obligations that EDSTA-enabling regulations are expected to impose on public-sector entities, including hospitals, and should inform a hospital's compliance-readiness work.

22 "Use of Artificial Intelligence in the Ontario Government: Independent Auditor's Report"

Article 3

The legal governance foundation

3.1 The governance statutory framework

Though AI-specific provisions are not specifically enumerated in the statutes that govern hospital directors' responsibilities in Ontario, the principles contained in the ONCA and the Hospital Management Regulation to the *Public Hospitals Act* (the Hospital Management Regulation) are broad enough to include AI governance in the exercise of directors' duties. The relevant sections of the two statutes are as follows:

- ONCA:[23]

  s. **21** Subject to this Act, the directors of a corporation shall manage or supervise the management of the activities and affairs of the corporation.

  ...

  s. **43(1)** Every director and officer of a not-for-profit corporation in exercising their powers and discharging their duties shall:

  (a) act honestly and in good faith with a view to the best interests of the corporation; and

  (b) exercise the care, diligence and skill that a reasonably prudent person would exercise in comparable circumstances.

- the Hospital Management Regulation:[24]

  s. 2(1) Every hospital shall be governed and managed by a board.

  ...

  s. **2(3)** The board shall,

  (a) monitor activities in the hospital for compliance with the Act, the regulations and the by-laws of the hospital;

  (b) take such measures as the board considers necessary to ensure that the provisions of the Act, the regulations and the by-laws of the hospital are being complied with; [...]

These two statutes form the statutory foundation for hospital directors' AI governance obligations. AI is an institutional activity at every hospital that has deployed it, and the board's obligation to govern and monitor extends to AI now, well before regulations are amended to specifically name it.

23 ONCA.

24 PHA.

ONCA imposes additional obligations for directors: compliance with the Act, regulations, and by-laws; a prohibition on waiving statutory duties; and a "reasonable diligence" defence where directors rely in good faith on professional reports and qualified management advice, provided they understand the matter sufficiently to supervise it.[25] The reliance defence does not make the board a passive recipient of management's narrative. The board retains the duty to ask informed questions and direct corrective action.

3.2 The standard of care: what "reasonably prudent" means in 2026

Canadian jurisprudence has established that directors will not be held to be in breach of the duty of care if they act "prudently and on a reasonably informed basis."[26] The decisions they make must be "reasonable business decisions in light of all the circumstances about which the directors or officers knew or *ought to have known*."[27]

This is the Canadian articulation of the "Business Judgment Rule" (BJR): a court will not second-guess a board decision made on an informed basis, in good faith, and through a defensible process.[28] The court looks to see that the directors made a reasonable decision, not a perfect decision, even if subsequent events cast doubt on the board's determination.[29] The BJR protects process and informed deliberation; not board inattention.

The BJR standard is objective, evolving, and has three direct implications for AI governance. First, lack of AI expertise is not a defence. Second, the standard includes what directors "ought to have known." A board that receives no AI governance reporting because it never asked for it has not met this standard. Third, the standard evolves with the environment. In June 2026, the reasonably prudent hospital director knows AI is deployed, understands the oversight obligations, receives structured reporting and directs corrective action when red flags arise. The spread of AI-enabled analytical tools may itself raise that standard, expanding what directors ought to have known and could have asked. Helleringer and Möslein describe this trajectory as an emerging AI judgment rule: as AI analytical and monitoring tools become more capable and accessible, a decision — or an oversight system — configured without their support may no longer be reasonably informed, which is the standard the BJR protects. This remains an academic argument about where the duty of care is heading, not established precedent or codified doctrine, and a board should treat it as a direction of travel rather than a present legal test. [30]

25 ONCA, s. 43(2) (duty to comply with Act and by-laws); s. 43(3) (duties non-waivable); s. 44 (reasonable diligence defence).

26 *Peoples Department Stores Inc. v. Wise*, 2004 SCC 68 at para 67 [*Peoples*].

27 Ibid.

28 See *Ford Motor Company of Canada, Ltd. v. Ontario Municipal Employees Retirement Board*, [2006] OJ No 27 at para 95, citing *Maple Leaf Foods Inc. v. Schneider Corp.*, (1998), 42 O.R. (3d) 177 (C.A.) at 192.

29 Ibid.

30 See "AI & the Business Judgment Rule: Heightened Information Duty", G. Helleringer and F. Möslein, University of Chicago Law Review Online (2025). See also "The Ungovernable Machine", Eran Kahana, Stanford Law School, CodeX, (March 17, 2026), characterizing the AI judgment rule as an emerging trajectory rather than established doctrine.

3.3 The duty to monitor: a distinctive hospital obligation

As mentioned in section 3.1, the Hospital Management Regulation imposes an express, freestanding duty to monitor institutional activities that is unique to Ontario hospital boards.[31] The principal Canadian corporate statutes, the *Business Corporations Act*,[32] the *Canada Business Corporations Act*,[33] the ONCA, and the *Canada Not-for-profit Corporations Act*, [34] do not use the word "monitor." Monitoring may be implied by the duty of care, but it is not expressly prescribed as an independent obligation applicable to directors. The Hospital Management Regulation changes that equation for every hospital board.

American corporate law has addressed the monitoring duty through two persuasive Delaware decisions:

- In *In re Caremark*, the court held that directors breach their duty when they utterly fail to implement any reporting system adequate to monitor material risk.[35] The court held that a board fulfils its monitoring obligation only by assuring itself "that the corporation's information and reporting system is in concept and design adequate to assure the board that appropriate information will come to its attention in a timely manner as a matter of ordinary operations."[36]

- In *Marchand v. Barnhill*, the Delaware Supreme Court extended that obligation to mission-critical operations, holding that where a risk is core to an organization's activities, a board with no monitoring mechanism cannot claim to have governed.[37]

These decisions are not binding in Ontario, but are analytically persuasive and align with Canadian jurisprudence on the fiduciary duty of directors in *Peoples Department Stores* and *BCE Inc. v. 1976 Debentureholders*. [38]

This monitoring duty has particular force for adaptive or self-modifying AI systems, including adaptive machine learning-enabled medical devices and agentic tools that can change their own behaviour after deployment. The seminal analysis of this exposure observes that each self-modification can overwrite the record linking an AI output to a traceable decision; so the reporting systems on which board monitoring depends may themselves be designed away unless the board requires otherwise.[39] The board need not understand the technical architecture. Consistent with its express monitoring duty under the Hospital Management Regulation, the board must satisfy itself that management

31 O. Reg. 965, s. 2(3).

32 *Business Corporations Act*, RSO 1990, c B.16.

33 *Canada Business Corporations Act*, RSC 1985, c C-44.

34 *Canada Not-for-profit Corporations Act*, SC 2009, c 23.

35 *In re Caremark International Inc. Derivative Litigation*, 698 A.2d 959 (Del. Ch. 1996) at 971 [*Caremark*].

36 Ibid at 970.

37 *Marchand v. Barnhill*, 212 A.3d 805 (Del. 2019) at 824 [*Marchand*].

38 Peoples at paras 63 and 67; BCE Inc. v. 1976 Debentureholders, 2008 SCC 69 at para 40 [BCE]

39 "The Ungovernable Machine", Eran Kahana, Stanford Law School, CodeX, (March 17, 2026). Kahana develops this analysis in respect of recursive self-improvement under the Delaware oversight doctrine in Caremark and Marchand. The principle applies with greater statutory force to the express monitoring duty under O. Reg. 965, and to adaptive or self-modifying AI deployed in hospitals.

maintains, and reports on, the change-control, logging and traceability, and human-approval safeguards that keep adaptive AI within its validated bounds — a function performed, for regulated devices, by the predetermined change control plan (PCCP) discussed in Article 6.

3.4 The proportionality principle: governance across the Ontario hospital spectrum

The AI governance oversight responsibilities are mandatory. Every hospital should have an AI governance policy, have named accountability, confirm PHIPA compliance, and verify Health Canada authorization for any SaMD that requires it. Proportionality governs the depth of implementation, not the existence of the obligation.

Court decisions outside Ontario have established different standards of care. For example, in *Bateman et al. v. Doiron et al*, the court held that the clinical standard of care applicable to a hospital is informed by institutional resources and context.[40] That principle does not translate to board governance obligations under the ONCA and the Hospital Management Regulation. In Ontario, the obligation to govern AI is not diminished by institutional size. Its implementation is proportionate to institutional complexity. Larger academic centres may deploy enterprise AI service models, while smaller hospitals opt for controlled AI pilot projects and shared-service adoption. Regardless of the form AI deployment takes, the governance obligations apply. Resource-constrained hospitals should engage with Ontario Health shared-service models and resist deploying AI they cannot govern safely.

3.5 Resource allocation: the decision to deploy AI

Hospitals have finite resources and numerous competing demands. Funding for AI initiatives competes for those finite resources in hospital budgets. The costs associated with AI deployment are complex, including licensing, integration, training, change management, and ongoing monitoring, in addition to acquisition alone. Furthermore, benefit realization at the institutional level is uncertain. The board will need clear cost, complexity, and expected-benefit picture for proposed deployment. A board that defers AI deployment, narrows its scope, or stages it incrementally is exercising the same fiduciary judgment as one that deploys AI broadly.

There are examples where courts have been deferential to boards' judgment calls on resource allocation. For example, in the recent decision *Abbott v. London Health Sciences Centre*, the Court of Appeal for Ontario affirmed that boards of public hospitals in Ontario are empowered to make decisions to manage hospital resources, including broad discretion to define the services the hospital provides pursuant to the PHA.[41] These principles have direct implications for the AI funding decisions of boards across Ontario.

The Supreme Court of Canada takes the same approach. Publicly funded health coverage is "by its very terms, a partial health plan" that does not promise to meet every medical need.[42] A reasoned, documented decision to fund AI at a given scale, including a decision to stage, narrow, or defer, is a defensible exercise of fiduciary judgment.

40 *Bateman et al. v. Doiron et al.* (1991), 118 N.B.R.(2d) 20 (Q.B.) at 14.

41 *Abbott v. London Health Sciences Centre*, 2025 ONCA 895 at para 64.

42 *Auton v. British Columbia (AG)*, 2004 SCC 78 at para 41; see also *Flora v. Ontario (HIP)*, 2008 ONCA 538 at paras 90-101.

Article 4

Legislation, regulations and regulatory guidance

In addition to the overarching governance legislation and jurisprudence in Article 3, boards must also consider legislation, regulations and regulatory guidance documents that are specific to the health sector. The following are the statutes, regulations and regulatory guidance that may be relevant to boards discharging their AI governance duties

- *Regulated Health Professions Act*, 1991 (RHPA)
- Cyber Security Regulation to the *Enhancing Digital Security and Trust Act* (the Cyber Security Regulation)[43], pursuant to the EDSTA

AI Scribe guidance issued by the Information and Privacy Commissioner of Ontario's (IPC) and the College of Physicians and Surgeons of Ontario's (CPSO).

4.1 The Regulated Health Professions Act framework: the diagnosis communication wall

Under Ontario's RHPA, communicating a diagnosis to a patient is a controlled act.[44] The RHPA prohibits any person from performing a controlled act when providing healthcare unless the act is performed by a regulated health professional authorized to perform the specific controlled act.[45] An AI system is not a regulated health professional and cannot legally communicate a diagnosis to a patient in Ontario, regardless of diagnostic accuracy.

In addition to developing and deploying policies and processes related to AI governance at the hospital-level, boards must be mindful of the obligations imposed upon regulated health professionals by the RHPA in areas such as:

- **Diagnostics**. Diagnostic tools may generate findings internally within the clinical team, but a regulated professional must communicate those findings to the patient. It is also critically important for clinicians to understand when the AI output generated from any diagnostic tool is intended to be used as a clinician decision support tool (i.e., one of many factors that a clinician should consider) vs. when AI output may be relied upon. See Article 6 for additional information on the distinction between clinical support decision tools and medical devices.
- **Patient-facing AI and virtual care**. Any AI tool that communicates directly with patients, such as virtual care chatbots, automated results platforms and AI symptom checkers, risk contravening

43 O. Reg. 51/26: Cyber Security, under EDSTA.

44 *Regulated Health Professions Act, 1991*, SO 1991, c 18, s. 27(2)1. "Communicating to the individual or his or her personal representative a diagnosis identifying a disease or disorder as the cause of symptoms of the individual in circumstances in which it is reasonably foreseeable that the individual or his or her personal representative will rely on the diagnosis." [*RHPA*].

45 RHPA.

the RHPA when they are not licensed SaMD operated in accordance with the product's intended and authorized use. Vendor contracts for patient-facing AI must confirm the system does not communicate diagnoses or otherwise perform controlled acts reserved for healthcare professionals directly without regulated professional review, unless the system is a licensed medical device. At the federal level, two bills introduced in June 2026 and not yet in force — Bill C-34, which would impose on operators of regulated "chatbot services" a duty to act responsibly, including crisis-intervention measures and safeguards against a chatbot deceiving users about its AI nature, and Bill C-36, which would repeal and replace Part 1 of the federal *Personal Information Protection and Electronic Documents Act*, recognize privacy as a fundamental right, and impose transparency and explanation obligations on automated decision systems — would, if enacted, bear directly on patient-facing AI of this kind.[46] In our detailed overview of Bill C-36, we canvass the bill's broader implications for organizational privacy governance, including its enforcement regime, automated decision system obligations, and new individual rights.

- **AI scribes and clinical documentation**. An AI scribe that generates clinical notes referencing a diagnosis does not violate the RHPA because the software is not communicating that diagnosis to the patient and the physician reviews, edits, and finalizes any draft output generated by AI prior to finalizing their clinical notes.

For each stream 2 deployment, hospitals must identify whether the tool requires human-in-the-loop oversight. High-impact, low-reversibility decisions such as diagnosis, treatment selection, medication dosing, or triage prioritization generally require human-in-the-loop oversight. The board must confirm that management has defined and enforced appropriate boundaries for every clinical AI deployment.

4.2 Cybersecurity regulation: three operational obligations

The first EDSTA regulation, the Cyber Security Regulation, takes effect July 1, 2026. It applies to every Ontario hospital graded Group A, B, or C under the PHA — broad categories that cover hospitals with fewer than 100 beds, hospitals with more than 100 beds and teaching hospitals and the University of Ottawa Heart Institute.[47] It imposes three obligations that must be operationalized now.

(a) Obligation 1: designated primary contact

Under s. 4 of the Cyber Security Regulation, every covered hospital must designate a senior management employee with decision-making authority over cybersecurity as its primary contact with the Ministry of Public and Business Service Delivery, and a second employee as alternate. Names, titles, and contact information must be provided to the Ministry's Chief Information Security Officer (CISO). Any change must be reported within 10 business days. The designated primary contact also approves the hospital's Cyber Security Maturity Assessment (CMA) summary before Ministry submission.

46 Bill C-34, *An Act to enact the Digital Safety Act and the Digital Safety Commission of Canada Act and to make consequential amendments to other Acts* (short title *Safe Social Media Act*), 1st Sess, 45th Parl, 2025–2026 (first reading June 10, 2026); Bill C-36, *An Act to enact the Protecting Privacy and Consumer Data Act, to amend the Personal Information Protection and Electronic Documents Act and to make amendments to other Acts*, 1st Sess, 45th Parl, 2025–2026 (first reading June 15, 2026). Both bills were at first reading as of the date of this guide and are not in force. See also "The Protecting Privacy and Consumer Data Act (Bill C-36): key obligations and enforcement overview", Osler, Hoskin & Harcourt LLP, (June 2026).

47 R.R.O. 1990, Reg. 964: Classification of Hospitals, under the PHA, s. 1(1); Cyber Security Regulation, s. 2.

(b) Obligation 2: cyber security maturity assessment

Every covered hospital must complete an initial CMA by July 1, 2027 and at minimum every two years thereafter.[48] Within 30 business days of each assessment, a summary must be submitted to the Ministry CISO including: methodology used; framework or model name; overall maturity score; and areas for improvement. The summary must be approved by the designated primary contact. A hospital that completed a qualifying CMA within one year before July 1, 2026 may treat it as the initial assessment.

(c) Obligation 3: critical cyber security incident reporting

Pursuant to s. 7 of the Cyber Security Regulation, every covered hospital must report a "critical cyber security incident" to the Ministry CISO within 72 hours of confirmation. A critical incident is defined as an incident that has impacted the "security, continuity, confidentiality, integrity or availability of digital information collected, used, retained or disclosed by the entity" or "the infrastructure housing or transmitting digital information" and that satisfies at least one of the following criteria: (a) significant adverse impact to public services (b) risk to public safety (c) significant recovery efforts or incident response plan activation or (d) significant risk to institutional reputation and public confidence.[49]

An AI system cyberattack of the type that occurred in Ascension Health, discussed in more detail at section 5.3, satisfies at minimum criteria (a) and (d) simultaneously.[50] The 72-hour report must include: primary contact and alternate names; dates and times; incident description and threshold justification; and description of information impacted.

This reporting obligation does not replace other potentially relevant reporting obligations: a single AI cyber incident may concurrently trigger PHIPA breach notifications.

The regulatory floor in the Cyber Security Regulation sits above an operational layer of federal and provincial cyber guidance directed at hospitals. The Canadian Centre for Cyber Security (CCCS), a federal agency that provides advice, guidance, services and support on cyber security for Canadians, issued the *Top 10 Artificial Intelligence Security Actions* (ITSAP.10.049), which sets out the operational framework for protecting against the adversarial use of AI, and the *Cyber Security for Connected Medical Devices* primer (ITSAP.00.132), which addresses the device-specific exposure layer that intersects with Health Canada-authorized SaMD.[51]

48 Cyber Security Regulation, ss. 5-6.

49 Cyber Security Regulation, s. 7.

50 "Angel Joel Gusman Negron et al v. Ascension Health et al.", United States District Court for the Eastern District of Missouri, (September 23, 2025). (Following a May 2024 cyberattack, a federal judge denied motion to dismiss, allowing negligence and negligence per se claims to proceed nationwide) [*Ascension Health*].

51 "Top 10 artificial intelligence security actions: A primer - ITSAP.10.049", Canadian Centre for Cyber Security. (March 2026). [*Top 10 AI Security Actions*]. "Cyber security for connected medical devices (ITSAP.00.132)" Canadian Centre for Cyber Security, (November 2021). [*Cyber Security for Medical Devices*].

The Ontario Health Cyber Security Centre, the cyber security division of the province's health care administration, released an advisory regarding phishing-resistant multi-factor authentication, least-privilege access, and centralized monitoring on May 1, 2026.[52]

Alignment with industry-specific guidance could support the Cyber Security Regulation compliance and inform the board's cyber-related materiality and reporting expectations.

4.3 IPC guidance

(a) IPC and Ontario Human Rights Commission *Principles for Responsible Use of AI*

In Ontario, the responsible use of AI is being contemplated by both regulatory and administrative bodies. The IPC and the Ontario Human Rights Commission have released six fundamental principles to guide the responsible use of AI in Ontario. The broad principles, released in January 2026, are intended to help organizations develop, deploy, and use AI in ways that maintain public trust by respecting privacy and human rights. Boards should consider the following high-level principles when developing an AI governance framework:[53]

- **Valid and reliable**. AI systems must exhibit valid, reliable, and accurate outputs.
- **Safe**. AI must be developed, acquired, and governed to prevent harm or unintended harmful outcomes that infringe upon human rights, including the right to privacy and non-discrimination.
- **Privacy protective**. AI should be developed using a privacy by design approach. Developers, providers, or users of AI systems should take proactive measures to protect the privacy and security of personal information and support the right of access to information.
- **Human rights affirming**. Protections must be built into the design of AI systems and procedures. Institutions using AI systems must prevent and remedy discrimination effectively.
- **Transparent**. Institutions that develop, provide, and use AI must ensure that these AI systems are visible, understandable, traceable, and explainable to others.
- **Accountable**. Institutions should implement a robust internal governance structure with clearly defined roles and oversight procedures, including a human-in-the-loop approach, to ensure accountability throughout the entire life cycle of their AI systems.

52 Ontario Health. "Threat Intelligence Advisory – Anthropic Claude Mythos". May 1, 2026. (Not available online) [*Ontario Health Threat Intelligence Advisory*].

53 "Principles for the responsible use of artificial intelligence", Information and Privacy Commissioner of Ontario and Ontario Human Rights Commission, (January 2026).

(b) IPC guidance on AI scribes

The *Personal Health Information Protection Act, 2004* (PHIPA)[54] is relevant to every AI system in an Ontario hospital processing personal health information.[55] The hospital remains the health information custodian regardless of the AI vendor. The IPC's guidance document *AI Scribes: Key Considerations for the Health Sector* (AI Scribes Guidance) defines PHIPA obligations for ambient AI documentation.[56]

Although the IPC AI Scribes Guidance includes detailed guidance for implementation, the following concepts are relevant for boards to be aware of:

- **Governance framework required**: Custodians should have an AI governance and accountability framework in place before introducing an AI scribe, including an AI governance committee, an AI risk management framework, comprehensive policies and procedures, and an ongoing monitoring and assessment process.
- **Privacy impact assessment (PIA) recommended**: A PIA should be conducted before deploying any AI scribe. PIAs should be reviewed and updated when the AI system changes or when a material change in data processing occurs.
- **Consent required**: There is no statutory provision in PHIPA that expressly permits collection of personal health information by an AI scribe without patient consent. Given that AI scribes operate by recording conversations with patients, express consent obtained before the encounter, recorded in the health record, is a mandatory requirement expected by health privacy regulatory authorities and multiple regulatory colleges across Canada, including the CPSO.
- **Human oversight recommended**: All AI-generated clinical records should be reviewed for accuracy before they are finalized, used or disclosed.
- **Considerations for vendor contracts**: The same principles that apply to any agent or electronic service provider also apply to AI-enabled technology. Among other things, vendor contracts should include: (a) prohibition on vendors from processing personal health information for any purpose unless authorized by the custodian in writing (b) mandatory notification when AI performance falls below accuracy thresholds and (c) an obligation to proactively notify the custodian when AI is producing unexpected outputs.
- **Ongoing monitoring and assessment**: The IPC guidance requires monitoring that is both regularly scheduled and triggered when unexpected behaviour thresholds are detected.

54 *Personal Health Information Protection Act, 2004*, SO 2004, c 3, Sch A [*PHIPA*].

55 This guide does not address all relevant privacy considerations applicable to hospital's engaging third party vendors and is limited to addressing certain limited AI considerations.

56 "AI Scribes: Key Considerations for the Health Sector", Information and Privacy Commissioner of Ontario. (January 28, 2026). [*IPC AI Scribe Guidance*]. See also "Artificial intelligence gone wrong: Why custodians need strong AI frameworks and policies", Information and Privacy Commissioner of Ontario, (April 27, 2026).

4.4 CPSO guidance

The CPSO's *Advice* regulator's expectations of physicians on the use of clinical AI is also relevant for boards to consider.[57] As noted above, the CPSO guidance requires express patient consent before ambient AI documentation tools are used in any clinical encounter.

The CPSO guidance highlights two key obligations:[58]

- **Informed consent for AI scribes**: Physicians recording conversations using AI must obtain express informed consent from patients before using the technology in the clinical encounter. This consent must be knowledgeable, voluntary, and recorded in the health record. This is a professional obligation for every Ontario physician deploying ambient AI scribes; therefore, every hospital with AI scribes must ensure a patient consent protocol meets this standard.

- **Clinical judgment**: Physicians exercise clinical judgment in connection with any AI-generated recommendation or output. Boards should be aware of the CPSO's guidance, including physicians' obligations to consider accuracy, accountability, data privacy and protection, potential for bias and transparency.[59]

Every hospital with an ambient AI scribe system must have a patient consent protocol that (a) provides express, knowledgeable patient consent before the AI scribe is used (b) allows patients to decline without affecting care quality and (c) records consent in the health record. The board should direct management to confirm a compliant protocol is in place for all current deployments.

57 "Advice to the Profession: Using Artificial Intelligence in Clinical Practice", College of Physicians and Surgeons of Ontario, (August 2025), [CPSO Advice].

58 CPSO Advice.

59 Ibid: "When integrating AI applications into their practice, physicians need to be aware of the following: (1) Accuracy: AI may sometimes produce false or inaccurate content. Physicians need to review all information generated using AI for accuracy and completeness (2) Accountability: physicians are ultimately accountable for their use of AI tools, including when using AI to support clinical decision-making or medical documentation (e.g., through AI scribes) (3) Data privacy and protection: all patient data entered into AI applications needs to be kept private and secure. Physicians' obligation to protect their patients' personal health information is no different when using AI than in any other circumstance (4) Potential for bias: AI can repeat biases in the data it learns from, which can have adverse impacts on certain groups. Physicians need to ensure the AI tool they are using is appropriate for their patient population, and evaluate the output of the tool for bias and (5) Transparency: physicians need to inform patients about how AI will be used, and in particular obtain patient consent before recording conversations using AI. Physicians are ultimately accountable for their use of AI tools, including when using AI to support clinical decision-making or medical documentation (e.g., through AI scribes).

4.5 Enforcement: what non-compliance costs

Hospitals that fail to comply with the legislative and regulatory requirements related to the deployment of AI tools may be liable for administrative monetary penalties or tort law damages.

Regulatory instrument	Enforcement exposure
PHIPA: IPC	Up to $100,000 for an individual; up to $500,000 for an organization. A PHIPA enforcement order is a public instrument.
EDSTA: Administrative monetary penalties	The penalty framework under EDSTA's enabling regulations is not yet finalized. The legislative architecture supports administrative monetary penalties for designated organizations that fail to implement required AI accountability frameworks.
Ontario tort law and *Public Hospitals Act*	A material AI-related patient harm event at an institution with no AI governance framework, no AI inventory, and no documented board oversight creates civil liability exposure and potential regulatory investigation by the Ministry of Health.

Article 5

Canadian jurisprudence: where the law stands and where it is going

5.1 Negligence: legal framework that will govern the first Canadian cases

As of June 2026, no Canadian court has apportioned liability specifically for AI-related patient harm in a healthcare setting. However health-related AI tools have been subject to regulatory scrutiny and an ongoing class action proceeding.[60] U.S. cases could be an indication of where Canadian jurisprudence could be heading.[61] The private law tort of negligence is a likely cause of action for individuals alleging harm from the improper deployment of clinical AI tools. The following are elements of the tort of negligence:

- **Duty of care**: Every Ontario hospital owes a private duty of care to its patients. That duty extends to the systems and tools the hospital deploys in patient care, including AI tools. As described above, the regulatory instruments defining the duty of care in the first Canadian cases are already in place.

- **Breach of the standard of care**: The plaintiff must prove the healthcare provider breached the standard of care. As AI tools become more widely adopted, the standard of care may evolve to require or expect the use of AI by clinicians in certain circumstances.

- **Causation**: The plaintiff must prove the breach of the standard of care, or statutory requirement, had a causal link to their injury. AI creates causation challenges because deep learning systems often cannot explain their reasoning. Where an AI tool contributed to patient harm, establishing the causal link requires evidence about training data, model architecture, and decision pathway, which may be invisible to the clinical team and which vendors may resist disclosing. The court-ordered discovery of algorithm development documents in *Lokken v. UnitedHealth Group* illustrates how causation may be contested in future Canadian court proceedings.[62]

A hospital can be held liable for systemic negligence, such as procurement failures, failure to train staff on AI tools, deployment of a known faulty system, or failure to monitor and, if appropriate, mitigate a deployed system.

60 See Section 5.2.

61 See *Ascension Health; Estate of Gene B. Lokken et al. v. UnitedHealth Group, Inc. et al.*, United States District Court for the District of Minnesota. 0:23-cv-03514. November 14, 2023 (A federal court allowed breach of contract and good faith claims to proceed) [*Lokken*]; *Raine v. OpenAI*, L.P. S.F. Superior Court, filed August 2025 (Negligence and wrongful death lawsuit alleging ChatGPT validated suicidal ideation for a 16-year-old) [*Raine*]; *Garcia v. Character Technologies, Inc. and Google LLC*. M.D. Fla., settled January 2026 (First resolved AI harm lawsuit involving a minor and an AI companion chatbot) [*Garcia*].

62 *Lokken*.

5.2 Proceedings related to AI use

Decisions by Canadian regulatory authorities from 2021 through 2026 have established the doctrinal framework that will govern legal disputes related to AI use. Each of the following is directly relevant to the governance obligations boards are discharging now.

Moffatt v. Air Canada.[63] The British Columbia Civil Resolution Tribunal held Air Canada liable for negligent misrepresentation based on inaccurate information from an AI chatbot on its website. Air Canada's argument that the chatbot was a separate legal entity was rejected. The tribunal held that Air Canada owed a duty of care to users of its chatbot and had failed to ensure the accuracy of the chatbot's representations. *Governance implication*: Any AI system communicating with patients (virtual care platforms, chatbots, AI-generated communications) is an exposure for which the hospital will be held accountable. The "it was the AI, not us" defence is not available.

Lam v. Flo Health Inc.[64] The Supreme Court of British Columbia certified a class action against an AI-based health tracking application for alleged privacy breaches involving personal health information. Certified claims include breach of privacy legislation, intrusion upon seclusion, and breach of confidence. *Governance implication*: AI-based health applications handling patient data face Canadian class action liability when data is handled inconsistently with privacy representations. Vendor contracts must accurately describe how the vendor processes patient data.

Clearview AI Inc. v. Information and Privacy Commissioner for British Columbia (2021-2026).[65] The B.C. Court of Appeal upheld B.C.'s provincial privacy order finding Clearview AI had contravened privacy legislation by scraping facial images without consent and confirmed that the "publicly available" exception must be narrowly construed given privacy's quasi-constitutional status. The Court modernized the "real and substantial connection" test: where an AI system's business model depends on large-scale online data collection and Canadian data is part of that collection, provincial privacy legislation applies regardless of vendor location. *Governance implication*: Vendor contract provisions warranting lawful training data provenance are the hospital's primary due diligence mechanism.

IPC Re Otter.ai.[66] The most directly relevant Canadian regulatory action on AI deployment in hospitals. A former physician's personal Otter.ai account joined a virtual hepatology rounds meeting, recorded the discussion of seven patients, and distributed the transcript to 65 invitees, including departed staff. The IPC's findings, immediate breach response protocols, approved AI tools list, firewall-level blocking, staff training, formal data deletion request, and a governance and privacy impact assessment, became the foundation for the IPC's AI Scribes Guidance. *Governance implication*: This is the shadow AI risk pattern Appendix 5 to Sample Board AI Governance Policy is designed to prevent.

63 *Moffatt v. Air Canada*, 2024 BCCRT 149, February 14, 2024.

64 *Lam v. Flo Health Inc.*, 2024 BCSC 391; Class action certification, 2025 BCSC 993.

65 *Clearview AI Inc. v. Information and Privacy Commissioner for British Columbia*, 2024 BCSC 2311; 2026 BCCA 67.

66 *Re Otter.ai*, Information and Privacy Commissioner of Ontario (File HR24-00691, published October 27, 2025) [*Re Otter.ai*].

Hospital for Sick Children v. Ontario (Information and Privacy Commissioner).[67] The Divisional Court upheld the IPC's finding that encryption of personal health information in a ransomware attack constitutes "use" and "loss" under PHIPA, triggering notification obligations even without evidence the threat actor viewed the information. This confirms that PHIPA's "use" and "loss" concepts will be interpreted broadly: any AI system interaction with PHI that could constitute "use" engages PHIPA obligations. It is also the leading authority on an Ontario hospital's direct accountability to the IPC. *Governance implication*: An AI-related PHIPA breach triggers IPC enforcement authority. The hospital's AI governance framework is the documentation produced in any IPC investigation.

5.3 Relevant U.S. Cases

Garcia v. Character Technologies / Google (M.D. Fla., October 2024; settled January 2026).[68] Megan Garcia's 14-year-old son died by suicide in February 2024 following prolonged interaction with a Character AI chatbot that had represented itself as a licensed psychotherapist. Garcia sued for negligence, wrongful death, and product liability, arguing that the companies failed to implement safety measures for vulnerable minors. In May 2025, the court rejected the defendants' motion to dismiss. The case settled in January 2026, the first resolved AI harm lawsuit involving a minor. Terms of settlement are not public.

Raine v. OpenAI (S.F. Superior Court, August 2025; ongoing).[69] The parents of 16-year-old Adam Raine allege that ChatGPT shifted from providing academic support to validating suicidal ideation, providing specific methods, and discouraging the teenager from telling his parents. The complaint alleges OpenAI launched GPT-4o after removing automatic termination protocols for conversations involving suicidal planning. This case is among a wave of similar actions across multiple U.S. jurisdictions.

Angel Joel Gusman Negron et al v. Ascension Health et al.[70] In May 2024, Ascension Health, operating some 140 hospitals across 19 states, suffered a ransomware attack that affected nearly 5.6 million patients. A class action followed within days and, in September 2025, a federal judge allowed nationwide negligence claims to proceed.[71]

Estate of Lokken v. UnitedHealth Group (D. Minn., 0:23-cv-03514).[72] Gene Lokken, 91, fractured his leg and ankle and was admitted to a nursing facility. UnitedHealth, using its nH Predict algorithm built on a database of six million patients, terminated his coverage. The algorithm is alleged to have had a 90% error rate. UnitedHealth knew this, but continued using it because only 0.2% of policyholders ever appealed. In February 2025, the federal court allowed breach of contract

67 *Hospital for Sick Children v. Ontario (Information and Privacy Commissioner)*, 2025 ONSC 5208 (Div. Court).

68 *Garcia*.

69 *Raine*.

70 Angel Joel Gusman Negron et al v. Ascension Health et al., United States District Court for the Eastern District of Missouri, (September 23, 2025).

71 Ibid.

72 *Estate of Gene B. Lokken et al. v. UnitedHealth Group, Inc. et al.*, United States District Court for the District of Minnesota. 0:23-cv-03514. November 14, 2023 (A federal court allowed breach of contract and good faith claims to proceed) [*Lokken*].

and implied good faith claims to proceed.[73] In March 2026, the court ordered production of the algorithm's development documents, training data, and source code. The governance lesson is that UnitedHealth's policy documents promised human clinical review and deploying an algorithm as the effective decision-maker, contrary to that representation, constitutes the breach.

Change Healthcare / The Wyden question (2024)[74]

Senator Ron Wyden's opening statement at the U.S. Senate Finance Committee hearing on the Change Healthcare cyberattack identified the absence of basic governance infrastructure as the central board failure. The Board director checklist in Schedule E is drawn from Wyden's framework, applied to AI governance. Every question Wyden asked about cyber is equally applicable to AI: Does the board know what systems are deployed? Does it receive structured reporting? Is there named accountability?

73 *Estate of Gene B. Lokken et al. v. UnitedHealth Group, Inc. et al* [PDF], United States District Court for the District of Minnesota. 0:23-cv-03514. (February 13, 2025).

74 Hacking America's Health Care: Assessing the Change Healthcare Cyber Attack and What's Next, Senator Ron Wyden, Opening Statement, United States Senate Finance Committee, (May 1, 2024).

Article 6

AI-enabled medical devices: what boards need to know

6.1 Overview

Hospitals must be aware of the legislative and regulatory framework governing medical devices in Canada, in particular Health Canada's regulation of SaMD. Out-of-the-box AI tools from credible, established vendors may simply require due diligence by hospitals to ensure they meet the statutory and regulatory requirements of medical devices.

Hospitals that are seeking to develop or commercialize proprietary AI tools or systems will require more robust board oversight and a risk mitigation framework to ensure compliance with applicable laws and the principles underlying the medical device authorization regime that are beyond the scope of this guide.

6.2 Health Canada: regulation of software-as-a-medical device (SaMD)

The regulatory foundation for medical devices in Canada is the *Food and Drugs Act* and the *Medical Devices Regulations*.[75] A "medical device" includes, among other things, any instrument, apparatus, contrivance or other similar article manufactured, sold, or represented for use in diagnosis, treatment, mitigation, or prevention of disease or abnormal physical state. Software that independently performs a medical function, including AI diagnostic tools, and medical imaging AI, as well as many other examples, are regulated as SaMD and are subject to the *Medical Devices Regulations*.[76]

When SaMD is classified as a class II, III or IV medical device, it must be authorized by Health Canada as a licensed medical device and issued a medical device licence before it is imported or sold in Canada. United States Food and Drug Administration (FDA) clearance is not interchangeable with a Health Canada medical device licence.

Manufacturers of AI-enabled SaMDs should consult Health Canada's Pre-market Guidance for Machine Learning-Enabled Medical Devices (Health Canada MLMD Guidance)[77], which applies to new and amendment licence applications for Class II, III and IV machine learning enabled medical devices and addresses the machine learning system across the device lifecycle (Class I to IV).

75 *Food and Drugs Act*, s. 31(1); Medical Devices Regulations.

76 Ibid.

77 "Pre-market guidance for machine learning-enabled medical devices", Health Canada, (April 2026).

The Health Canada MLMD Guidance outlines the technical, lifecycle and post-market monitoring information that manufacturers should consider in demonstrating the safety and effectiveness of AI/ML-based SaMD, applying a risk-based approach rather than prescribing fixed requirements for all devices.[78]

Core elements addressed in the Health Canada MLMD Guidance include

- abide by "good machine learning practice" when designing, developing, evaluating, deploying and maintaining an MLMD
- use data to develop an MLMD that is representative of the Canadian population and clinical practice, by applying sex- and gender-based analysis plus (SGBA Plus) across the device lifecycle and using disaggregated subgroup data to promote equitable and representative performance across Canada's diverse population
- provide appropriate clinical evidence, including clinical validation studies, to support the safe and effective clinical use of the device (expected for Class III and IV applications, and available upon request for Class II)
- ensure transparency by communicating clear information about the device's development, performance and limitations to patients, users, health care providers and regulators, including through device labeling and the software user interface
- consider whether a predetermined change control plan (PCCP) will be included in the MLMD licence application to provide a mechanism for Health Canada to address cases where the regulatory pre-authorization of planned changes to machine learning systems is needed to address a known risk
- anticipate that Health Canada may impose or amend licence terms and conditions under subsection 36(2) of the Medical Devices Regulations to manage uncertainties relating to machine learning systems and PCCPs across the device lifecycle
- continue to observe the safety and effectiveness requirements applicable to MLMDs following the pre-market phase throughout the entire product lifecycle
- develop post-market monitoring plans and include in MLMD licence application, a description of the processes, surveillance and performance monitoring plans, and risk mitigation strategies to ensure ongoing performance

For AI products considered software as a medical device, medical device manufacturers are required to manage device-level risk across the product lifecycle in a manner consistent with the recognized international standard for medical-device risk management, ISO 14971[79], and with the

78 Health Canada published medical device guidance documents in March and April 2026 clarifying certain lifecycle regulatory expectations for Class II to IV devices, including terms and conditions; the interpretation of "significant change"; clinical evidence requirements; pre-market guidance for machine learning-enabled medical devices; and summary reports and issue-related analyses. In particular, the Health Canada significant change guidance clarifies the threshold for a "significant change" requiring a licence amendment for Class III and IV devices, particularly for software, performance, cybersecurity, and compatibility changes.

79 ISO 14971:2019, *Medical devices — Application of risk management to medical devices*.

risk-management expectations in the Health Canada MLMD Guidance — identifying erroneous outputs, bias, model-fitting failures, and performance degradation, applying risk controls, and documenting residual-risk acceptability. As a result, ensuring that AI-enabled software that is considered a Class II, III or IV medical device has been issued a medical device licence by Health Canada. Understanding the medical device licence's terms and conditions, if applicable, will be of critical importance, given that hospitals will not be in a position to assess whether vendors of software products have satisfied all applicable requirements Health Canada assesses prior to issuing a medical device licence.

At a board level, every Ontario hospital deploying AI classified as a Class II, III or IV SaMD should have procedures in place to ensure that hospital staff assess whether vendors of SaMD products hold valid Health Canada authorization. When a SaMD vendor is outside Canada, the hospital could be considered the importer of the SaMD, in which case the hospital could become directly liable for ensuring the SaMD has all necessary Health Canada authorizations. Importing an AI-enabled SaMD that requires a medical device licence that has not been issued a medical device licence by Health Canada is an offence under the *Food and Drugs Act*, which has a maximum fine of $5 million per offence.[80]

It may also be the case that certain software that appears to provide a medical purpose could be outside the scope of the Medical Devices Regulations. For example, Health Canada provides four exclusion criteria applicable to patient decision support software and clinical decision support software and, if all four are satisfied, the software falls outside regulation.

Critically, it is only when all four of the exclusion criteria below are satisfied that software for patient decision-making or clinical decision-making would not be considered a medical device:

1. Software that is not intended to acquire, process, or analyze a medical image or a signal from an in vitro diagnostic device or a pattern/signal from a signal acquisition system
2. Software that is intended to display, analyze, or print medical information about a patient or other medical information (such as demographic information, drug labelling, clinical guidelines, studies, or recommendations)
3. Software that is only intended to support a healthcare professional, patient or non-healthcare professional caregiver in making decisions about prevention, diagnosis, or treatment of a disease or condition, and
4. Software that is not intended to replace the clinical judgment of a healthcare professional to make a clinical diagnosis or treatment decision regarding an individual patient.

80 *Food and Drugs Act*, s. 31.2(1).

6.3 The director shield and its limits: PHA Section 13 and federal regulatory exposure

Section 13 of the PHA contains an extraordinarily broad shield for hospital directors.[81] For most governance decisions, a director acting in good faith under the PHA is protected from personal liability. However, PHA s. 13 is provincial and cannot shield directors against federal regulatory enforcement, including Health Canada enforcement under the *Food and Drugs Act* and the *Medical Devices Regulations*.[82] A hospital acting as an importer of a class II, III or IV AI SaMD without valid Health Canada authorization could potentially face enforcement action from Health Canada. There is also the potential for directors and officers to face personal liability for offences under the *Food and Drugs Act* and section 13 of the PHA provides no protection in this context.

This is one of the very rare situations in which an Ontario hospital director could potentially face unshielded personal liability. The combination of PHA s. 13 and the BJR provides strong protection for most governance failures that a board may commit, but falls short in connection with violations of federal legislation. As a result, the board should ensure that hospital staff verify and document Health Canada authorization before procuring or deploying any clinical AI qualifying as a medical device obtained from a third party vendor, whether or not the third party vendor is charging the hospital a fee for use of the product.

81 PHA, s. 13(1): "No action or other proceeding for damages or otherwise shall be instituted against any member of a... board... for any act done in good faith in the execution or intended execution of any duty or authority under this Act or the regulations or for any alleged neglect or default in the execution in good faith of any such duty or authority."

82 *Food and Drugs Act*, s. 31(1) (personal liability of directors and officers); Medical Devices Regulations, SOR/98-282.

Article 7

Sample governance architecture

7.1 The five-layer governance architecture

To govern AI safely and lawfully in Ontario, boards should anchor oversight in binding legal requirements, classify and manage AI risks within existing enterprise risk structures, and ensure clear accountability from the boardroom to the bedside. This guide sets out a five-layer governance architecture for consideration that begins with the legal and regulatory floor, builds tiered organizational oversight, and operationalizes deployment through a disciplined stage-gate process with defined artefacts, owners, and decision records. It concludes with continuous monitoring and incident response aligned to provincial cyber obligations and health-system guidance. Section 7.11 addresses in detail how this governance architecture integrates with the hospital's defences against AI-enabled cyber threats.

(a) Layer 1: legal and regulatory floor

No governance framework is adequate if it does not address operative legal obligations, such as the Cyber Security Regulations under EDSTA. The relevant statutes, regulations and guidance documents that form this non-negotiable foundation as of the date of this guide are set out in Article 4 and Article 6, but governance frameworks should be reviewed at a regular cadence to ensure that any new applicable legislation has been taken into consideration.

(b) Layer 2: risk identification and classification

Every AI tool should be inventoried and assigned a risk tier, and all AI risks should be integrated into the hospital's existing Integrated Risk Management (IRM) framework. In our view, AI should not be considered a new IRM domain. It should, instead, be categorized as a risk sub-category within each existing domain, surfacing within patient safety, digital health and cybersecurity, legal and regulatory, financial, and human resources. This approach ensures AI risks are governed within existing processes.

Within the Digital Health and Cybersecurity IRM domain, personal-device access to the EMR and hospital meeting platforms warrants particular attention as a discrete, AI-amplified risk sub-category.[83] Personal device access is also the surface on which the Ontario Health Cyber Security Centre's May 1, 2026 advisory directs phishing-resistant multi-factor authentication, least-privilege access, and centralized monitoring.[84] Hospitals should maintain a current policy governing personal-device access to the EMR and to hospital meeting platforms, reviewed by the AMC, where applicable.

(c) Layer 3: tiered organizational oversight

Governance capacity should be built at every organizational tier, with each layer requiring greater depth than the one above. The board oversees AI strategy, and the C-suite deploys AI initiatives and is accountable for compliance and performance. Finally, the operational and clinical tier

83 This risk is illustrated in *Re Otter.ai*.

84 Ontario Health Threat Intelligence Advisory.

executes AI use safely at the point of care. Board members and senior executives must maintain sufficient AI fluency to discharge their oversight duties, while frontline leaders must develop the competency to manage day-to-day deployment risks.

(d) Layer 4: operational and clinical governance

This level is where AI governance meets deployment. Governance of each AI deployment follows a stage-gate discipline

- intake, where the use case and owner are defined
- triage, in which risk tier is assigned and routing determined
- assessment, in which risk questions are answered and evidence gathered, including the Cyber Security Regulation cyber security exposure for tools with significant EMR access or network integration
- build/Buy, in which vendor or design are reviewed, including contractual cyber security obligations and incident notification requirements
- validation, in which the local test plan is executed, including security validation
- approval, in which conditions are documented and sign-off is obtained, including cyber security conditions for Tier 2 AI risk and above
- monitoring, in which incidents and drift are tracked continuously and the Cyber Security Regulation critical incident pathway activated if applicable
- retire, in which sunset criteria are defined and applied

Every stage should produce a document, an owner, and a decision record. Records of each decision could include:

1. Intake form: defines the problem, owner, intended users, data inputs, expected benefits, and measurable success criteria before any work begins
2. Rapid risk triage: assigns low, medium, or high risk and routes the use case to the appropriate approval pathway
3. Risk assessment and conditions: deepens the review for medium and high-risk use cases, documents sign-offs, and records the conditions attached to approval
4. Minimum viable product control checklist: confirms the minimum baseline controls are in place for delegated low-risk approvals
5. Risk Register entry: records material risks, mitigations, owners, and residual risk for ongoing tracking, and
6. AI inventory entry: the system of record for lifecycle status, conditions, monitoring, incidents, and next review date.

(e) Layer 5: monitoring and incident response

Governance does not end at deployment. Post-deployment clinical AI monitoring should be mandatory, automated, and continuous. Four post-go-live failure modes account for the majority of AI governance failures in practice and should be addressed as discrete sub-categories:

1. vendor model changes without notification: the deployed algorithm is silently updated, invalidating local validation
2. workflow drift: the clinical workflow in which the AI operates changes after approval, placing the AI outside its validated context
3. weak monitoring ownership: no named individual is accountable for ongoing performance tracking after the approval committee has discharged its pre-deployment role, and
4. use case scope creep: a tool approved for a specific clinical purpose is quietly applied beyond its original scope without re-assessment or re-approval

Each failure mode requires a named owner, a defined detection mechanism, and a defined response. Monitoring plans must address three metric types: outcome metrics that evaluate whether the AI is achieving its objective; process metrics that assess whether the AI is being used as intended; and balancing metrics that evaluate whether the AI is creating unintended adverse effects. Clinician overrides should be tracked as a primary monitoring signal. Consistent overrides by clinicians indicate either a performance problem or a training problem, and the framework must define a response to either issue.

7.2 The statutory architecture: board, CEO, Chief of Staff, and board committees

An Ontario public hospital board has two direct employees: the CEO (the "administrator" within the meaning of the PHA and the Hospital Management Regulation), and the Chief of Staff. Each reports to the board on distinct domains. The CEO is accountable for hospital operations and regulatory compliance. The Chief of Staff chairs the MAC[85] and reports to the board[86] on individual physician performance and the practice of medicine in the hospital.

The Quality Committee, established under the *Excellent Care for All Act, 2010* (ECFA)[87], oversees systemic quality.[88] The MAC, chaired by the Chief of Staff, reports to the board on the practice of medicine in the hospital under the Hospital Management Regulation.

Under the Hospital Management Regulation, the CEO is operationally accountable to the board for compliance with the PHA, the regulations, and the by-laws[89]. AI deployment is within that accountability perimeter.

85 Section 2(3)(c) of the Hospital Management Regulation.

86 Section 7(5) of the Hospital Management Regulation.

87 s. 3(1).

88 *Excellent Care for All Act*, SO 2010, c 14 [*ECFA*].

89 s. 3.

In line with these statutory responsibilities, the Audit/Finance Committee could have oversight over stream 1 (operational efficiency AI) and the Quality Committee could have oversight over stream 2 (clinical AI). See Sections 7.7, 7.8, 7.9 and 7.10 for additional considerations relevant to clinical AI oversight. The CEO supports both the Audit/Finance and Quality Committees in their AI oversight functions. Hospitals could consider implementing an AMC, reporting to the CEO, who carries AMC recommendations to the relevant stream committee on the materiality thresholds set out in the board's AI governance policy.

The statutory architecture is summarized below.

Role	Function / reporting structure	Statutory authority
Board	Governs and manages the hospital; monitors compliance with the Act, regulations, and by-laws; directors must demonstrate AI competency as outlined in Section 7.5.	ONCA s. 21; s. 43; PHA; O. Reg. 965, s. 2(3).
CEO (Administrator)	Responsible for compliance with the PHA, regulations, and by-laws. Operationally accountable for AI deployment. Reports to board (through Audit/Finance and Quality Committees).	O. Reg. 965, s. 3.
Chief of Staff / MAC Chair	Chairs the MAC; ex officio board member under hospital by-laws; reports on activities of the MAC, individual physician performance, and the practice of medicine. Distinct from VP Medical and CMIO, who are CEO-reporting employees. Reports to board (directly).	O. Reg. 965, s. 2(3)(c); PHA s. 35; O. Reg. 965, s. 7.
Audit / Finance Committee	CEO, or designate, is supporting senior employee. Reports to board.	Board by-law.
Quality Committee	Independent statutory monitor of systemic quality; chaired by a voting director. Reports to board.	ECFA s. 3; O. Reg. 445/10.[90]

90 O. Reg. 445/10: General, under the ECFA.

Role	Function / reporting structure	Statutory authority
Medical advisory committee	Individual physician performance: privileges, credentialing, supervision, conduct. Systemic quality recommendations routed to Quality Committee. Reports to board on individual matters and Quality Committee on systemic matters.	PHA s. 35; O. Reg. 965, s. 7.
AI management committee	Recommended senior management body executing AI deployment governance, supporting the Quality Committee, who could be responsible for stream 2, the Audit/Finance Committee, who could be responsible for stream 1, and the MAC on clinical AI; integrates the Cyber Security Regulation obligations. See Appendix 6 to the Sample Board AI Governance Policy (Schedule D). Reports to CEO.	Not applicable. Recommended for consideration under this guide.

7.3 AI risk tier classification and approval authority

The following table includes governance obligations that boards could consider implementing, based on an AI risk tier classification, with Tier 1 AI risk representing the lowest risk, and risk increasing incrementally through to Tier 4 AI risk. The AI risk tier classifications associated with each type of AI deployment set out in this guide have been adapted from the principles of the *European Union Artificial Intelligence Act* and the National Institute of Standards and Technology's Artificial Intelligence Risk Management Framework for use in the health care sector.[91]

In this guide, we have proposed a governance model that requires every AI system to be classified and approved before deployment, with approval authority escalating as the risk profile increases, with corresponding board reporting obligations. The approval authorities range from CIO/CTO for Tier 1 AI risk operational efficiency AI tools to board approval for agentic AI in accordance with the following:

91 "Artificial Intelligence Risk Management Framework" [PDF], National Institute of Standards and Technology, (January 2023).

Risk Tier	Example classification	Example approval authority	Example board reporting obligation
Tier 1	Operational efficiency AI (For example, patient flow optimization, surgical scheduling, bed allocation, staff scheduling, supply chain, billing workflows, and other operational functions without direct clinical decision-making)	CIO/CTO	Annual AI performance and compliance report
Tier 2	Clinical decision support AI (For example, deterioration monitoring, readmission risk, diagnostic flagging without autonomous output)	AMC recommendation; stream 2: senior medical officer endorsement and CEO approval; CFO concurrence	Semi-annual AI governance executive report; material events within 48 hours
Tier 3	Autonomous diagnostic AI / SaMD (For example, medical imaging AI, AI generating autonomous clinical recommendations, regulated AI-enabled medical devices)	Senior medical officer endorsement and CEO approval; privacy officer / legal counsel review required	Semi-annual AI governance executive report; material events within 48 hours
Tier 4	Agentic AI / high-risk SaMD (For example, AI autonomously executing multi-step workflows across clinical systems without per-step human authorization)	Board approval required	Prior board approval required; material events within 48 hours

The Quality Committee's role in clinical AI is oversight, not per-deployment approval: management approves stream 2 deployments that meet the guardrails in the board's AI governance policy, and the Quality Committee monitors systemic quality through the AMC's structured reporting and the

escalation of any matter exceeding the policy's materiality thresholds. Where the senior medical officer declines to endorse a Tier 3 deployment, it may not be approved by the CEO without the Quality Committee's endorsement.

To implement this governance framework, a material change to a deployed tool's autonomy, intended use, or underlying model may require re-classification and re-approval at the revised tier, and the tier classification and approval pathway of each AI tool should be documented in the AI inventory before deployment. Real-time intraoperative diagnostic AI, of the kind discussed in the University of Waterloo and Princess Margaret/University Health Network case study at Schedule B, illustrates the upper end of Tier 3, since it is an autonomous diagnostic AI / SaMD generating outputs that influence irreversible surgical decisions in real time. Though it has serious clinical implications, it does not require a Tier 4 AI risk rating, but Tier 3 AI risk governance applied at maximum rigour, including documented intraoperative fall-back protocols, demographic-subgroup performance evidence and, when applicable, Health Canada SaMD authorization specific to the intended intraoperative use case.

7.4 AI risk sub-categories within existing IRM domains

As discussed in section 7.1, every Ontario hospital's IRM framework should be revised to add AI, where appropriate, as a sub-category of existing IRM domains. As of the issuance of this guide, Healthcare Insurance Reciprocal of Canada's (HIROC) IRM taxonomy did not identify artificial intelligence as a discrete governance risk, and Accreditation Canada's Qmentum standards do not name AI as a governance standard.[92] However, we have provided examples of AI subcategory risks in the table below as illustrative examples. These examples do not represent an exhaustive list of all AI subcategories.

92 "2025 Top Healthcare Risks Report (IRM)", HIROC, (August 2025).

IRM domain	AI risk sub-categories that could be added
Patient safety	Hallucination producing incorrect clinical outputs; automation bias causing over-reliance on erroneous AI recommendations; diagnostic AI performing inconsistently across demographic groups; algorithm drift as clinical guidelines evolve; delayed or failed incident reporting for AI-related adverse events; communication risk: AI systems are probabilistic and cannot give precise recommendations; where clinicians do not understand how a recommendation was generated, the risk of misuse and alarm fatigue increases; false positive and false negative thresholds must be defined before deployment and monitored after it. Six AI-specific bias categories could be assessed (1) historical bias: data that has evolved without flagging the distinction (2) representation bias: datasets that exclude relevant patient subgroups (3) measurement bias: inconsistent data collection (4) aggregation bias: unique subgroup characteristics lost in combined data (5) evaluation bias: training data different from deployment context (6) deployment bias: system used beyond its approved scope. Model-fitting failures such as overfitting and underfitting can also produce outputs that do not generalize to the deployment population.
Care quality	AI clinical decision support not reflecting current clinical evidence; ambient AI documentation misrepresenting the clinical encounter without being corrected by clinicians; performance degradation in tools not re-validated after algorithm updates; AI scribe consent failure; explainability failure: AI systems that cannot explain how they generated a recommendation cannot be challenged by clinicians, cannot be investigated after an adverse event, and do not support the independent clinical judgment the RHPA and CPSO require.
Post-deployment monitoring failure (a cross-cutting failure mode spanning the domains above, rather than a discrete IRM domain)	Four post-go-live failure modes described in layer 5 of the five-layer governance architecture (section 7.1) govern many AI governance breakdowns. Additional failures could include no defined thresholds for retraining or retirement; no demographic subgroup performance tracking; and algorithm drift undetected as clinical guidelines evolve. For adaptive or self-modifying AI, additional Health Canada-recognized sub-categories include machine-learning performance degradation over time, post-deployment shifts in the input data distribution, and model changes implemented outside a PCCP.

IRM domain	AI risk sub-categories that could be added
Digital health and cybersecurity	AI systems with broad EMR access expand the cybersecurity attack surface; agentic AI creates the largest new attack vectors; vendor AI system cyberattack (see *Ascension Health* at Section 5.3). Shadow AI: employee use of personal consumer AI tools (ChatGPT, Gemini, Copilot) for hospital work without institutional authorization creates simultaneous PHIPA, CPSO, and EDSTA exposure and is the most prevalent uncontrolled AI risk at hospitals today. Sub-risks: clinicians entering patient data into unapproved LLMs; administrative staff processing confidential data outside approved platforms; unapproved tools supplementing diagnostic workflows. Most shadow AI reflects gaps in approved tooling and AI literacy.[93] Additional required sub-categories: traceability and auditability obligations: AI systems must maintain clear audit trails that can support incident investigation and regulatory inquiry; agentic AI misalignment (AI agents that plan and act autonomously create attack vectors and accountability gaps that exceed those of decision-support AI); data poisoning: adversarial manipulation of training or input data to alter AI system outputs (research has demonstrated that modifying even a single pixel in an image ingested by a diagnostic AI can materially alter its classification, making data integrity protection during inference; not just at training, a discrete security requirement); deepfake and synthetic content threats: AI-generated clinical documentation, images or communications designed to deceive clinicians or administrators; business continuity gaps if the hospital's business continuity plan does not address AI system outages as a critical operational dependency; not merely traditional IT failures.

93 See "Use of Artificial Intelligence in the Ontario Government: Independent Auditor's Report". The Office of the Auditor General of Ontario found that between April 2025 and August 2025, 12,000 Ontario public service staff accessed approximately 400 AI-related websites. Of those, about 60% were deemed unsafe or unsecured. The report also stated that the Ministry of Public and Business Service Delivery and Procurement had not implemented security controls to prevent Ontario public service staff from inadvertently uploading Ontarians' personal information, such as health cards, driver's licences and credit card information, or sensitive corporate data, to these AI websites.

IRM domain	AI risk sub-categories that could be added
Legal, regulatory and compliance	SaMD requiring a Health Canada licence deployed without Health Canada authorization; PHIPA non-compliance in AI data processing; EDSTA accountability obligations unmet; IPC and CPSO AI scribe guidance requirements unaddressed; OHIP billing anomalies from AI documentation; and shadow AI: unauthorized employee AI use, simultaneously constitutes a PHIPA breach, CPSO consent violation, and EDSTA accountability failure; adaptive MLMD modified outside an authorized PCCP, or a significant change to a Class III or Class IV medical device implemented without a Health Canada licence amendment. See Digital health and Cybersecurity domain and Appendix 5 to Sample Board AI Governance Policy.
Financial	OHIP coding compliance exposure; AI tool costs exceeding projected value; vendor lock-in reducing strategic flexibility; class action liability from AI-related patient harm.
Human resources	Clinician adoption resistance; de-skilling through over-reliance on AI clinical decision support; inadequate automation bias training; inadequate frontline IT staff training on AI system selection, implementation, maintenance and inspection. For academic hospitals: trainee automation dependency; concerns about AI impact on academic recognition; failure to comply with the *Working for Workers Four Act, 2024*[94] amending the *Employment Standards Act*, requiring Ontario employers to disclose in publicly-advertised job postings whether AI was used to screen, assess, or select applicants.
External relations and reputation	Public trust erosion from AI errors; media exposure from an AI-related patient harm event; published IPC enforcement order. For academic hospitals: reputational risk from research AI misuse significantly amplified.

The risks catalogued above are examples of why the opportunities discussed herein require governance *before* deployment; not instead of it. A board that builds the architecture to manage these risks simultaneously builds the architecture to enable AI's benefits.

94 *Working for Workers Four Act, 2024*, SO 2024, c 3.

7.5 AI competency architecture

Governance does not require AI expertise on the board. It requires AI fluency across the institution, built through structured orientation, ongoing education, and monitoring and reporting that surfaces gaps before they become adverse events. Each tier has both a fluency obligation and an operational accountability obligation.

(a) Board level

AI fluency at the board level could be built through:

- targeted recruitment identifying AI literacy and cybersecurity governance as sought competencies
- structured orientation ensuring every new director understands the AI inventory and regulatory obligations
- a rolling two-year education curriculum tied to regulatory developments, emerging case law, and the institution's deployment profile

Survey evidence associates a critical mass of three or more AI-literate directors with stronger oversight, reinforcing the targeted-recruitment objective above without making board-level AI expertise a precondition.[95] A single annual briefing on AI would not meet the expected standards. Where board expertise is absent, an external advisor with a direct reporting line to the relevant committee could be implemented as the compensating mechanism.

(b) C-suite tier

The board holds the CEO accountable for C-suite AI fluency. Each executive with AI governance responsibility must have working knowledge sufficient to make informed deployment decisions, identify governance gaps before they become adverse events, and provide independent assessments rather than progress narratives.

Operationally, the CEO could designate an executive owner for each AI context. For example, the senior medical officer could be the executive owner of stream 2 (clinical AI) and the CIO/CTO for stream 1 (operational efficiency AI). These owners could be responsible for establishing the AMC, maintaining the AI inventory, documenting Health Canada medical device licence verification for qualifying SaMD, completing PHIPA PIAs for AI systems processing patient data, and ensuring material vendor contracts have been appropriately reviewed.

The CEO should report to the board at a regular established cadence with material incidents reported to the board within 48 hours.

95 "Canada's AI Adoption Gap", Hartmann & Hartmann, (January 2026).

7.6 Operational Tier

Each hospital should develop and maintain formal protocols for the appropriate selection, implementation, training, maintenance and inspection of AI products, ensuring all staff interacting with AI systems are appropriately trained for their specific roles. To be effective, protocols should address:

- role-specific training calibrated to the tool's risk tier and the staff member's operational responsibilities[96]
- mandatory automation bias awareness training for all clinical staff using AI decision support tools
- documentation of training completion as a precondition to system access for higher risk deployments
- mandatory retraining when AI systems are materially updated or replaced
- regular competency review linked to post-deployment monitoring findings

(a) Training

Failure to train staff on deployed AI tools is a documented institutional liability exposure. Training must, in plain language, address not only how to use the AI tool, but how it generates its recommendations. Hospitals should establish a feedback loop for reporting incorrect AI outputs, with a documented process for escalating to the AMC and the applicable vendor. Each deployment should define in advance how conflicts between AI recommendations and clinician judgment are resolved, where applicable. The clinician should always make the final decision when providing healthcare services whether to rely on the output from any SaMD, which must be explicit in protocols.

(b) Monitoring and reporting as the enforcement mechanism

The monitoring system at the operational tier makes the governance architecture operational. The Sample Board AI Governance Policy provides a sample of some of the types of mechanisms that could be implemented, such as: template-driven semi-annual reporting (Appendix 1 and Appendix 2 to the Sample Board AI Governance Policy); defined thresholds triggering board review; a 48-hour material incident protocol (Appendix 4); and continuous post-deployment monitoring for all Tier 2-4 AI risk tools. It is also recommended that every material AI governance discussion be recorded in AMC meeting minutes.

The board's periodic AI reporting should be comprehensive. A sample reporting pack spans five domains set out in Appendix 2 to Sample Board AI Governance Policy. Typically immediate escalation is required when a patient safety, privacy, cyber, or compliance event creates material exposure; a vendor shift alters the risk profile; a key control fails; or a major capital decision falls outside approved guardrails.

96 See "Use of Artificial Intelligence in the Ontario Government: Independent Auditor's Report". In the Ontario public sector context, the Office of the Auditor General of Ontario is recommending that the Ministry of Public and Business Service Delivery and Procurement implement measures to ensure that all staff take the training on AI-related risks with refresher training offered as the Ontario public service advances AI use. Only three percent of Ontario public service staff completed the voluntary AI training course, Responsible Use of AI, between when it was launched in January 2024 and August 2025.

7.7 Quality Committee oversight of clinical AI deployment

Clinical AI deployment could be overseen at the board level by the Quality Committee. The Quality Committee could receive the AMC's reports and recommendations regarding stream 2 deployments through the CEO, and the MAC's reports and recommendations regarding systemic or recurring quality of care issues directly under the Hospital Management Regulation, s. 7(7).

7.8 The medical advisory committee's clinical input

The MAC is the principal locus of organized physician leadership in the hospital. The MAC is chaired by the chief of staff, an ex officio member of the board under the hospital's by-laws and the principal channel through which the medical staff's clinical judgment reaches the board. Its membership comprises the chiefs of department and elected representatives of the medical staff, with the senior medical officer, the chief nursing executive, and other senior employed leaders attending ex officio in a non-voting capacity. The MAC's perspective will be important for the governance of clinical AI deployment.

The MAC could be consulted on proposed clinical AI deployments through the chief of staff/MAC chair, with operational support from the senior medical officer and, in larger institutions, a CMIO. The MAC, supported by that team and by the AMC, could consider and provide recommendations on clinical use, workflow integration, rules of practice affecting medical staff, automation-bias and consent considerations, demographic-subgroup performance concerns, vendor model-update implications, and any other concerns the MAC identifies. Those recommendations could flow back through the chief of staff/MAC chair to the board on matters within s. 7(5) of the Hospital Management Regulation, and through the CEO to the Quality Committee on systemic matters.

Traditional clinical risks are individual and event-based, such as a single physician's judgment in a single patient encounter. AI risks operate at scale and propagate horizontally. The MAC is structurally well-suited to surface individual clinician concerns. The Quality Committee is structurally responsible for systemic quality. Clinical AI requires both, working together. Neither operating alone would be sufficient.

The AMC could support the MAC; not the other way around. The chief of staff/MAC chair, the senior medical officer, and (in larger institutions) the CMIO could be the conduit between the MAC and the AMC. The AMC could deliver structured assessments, including regulatory authorization status, privacy impact, vendor performance data, demographic-subgroup analysis, post-deployment monitoring evidence, automation-bias mitigation, and the use-case scope conditions, to the MAC ahead of any MAC consideration of clinical AI. The MAC could test those assessments against clinical reality. Where the MAC identifies a systemic issue, the issue could be routed to the Quality Committee under the Hospital Management Regulation, s. 7(7) and, where individual physician practice is implicated, the matter is handled under s. 7(5). The same incident could engage both pathways.

7.9 Systemic quality issues

Where the MAC identifies a systemic or recurring quality of care issue, the Hospital Management Regulation directs the reporting flow:

> **s. 7(7)** Where the medical advisory committee identifies systemic or recurring quality of care issues in making its recommendations to the board under subclause (2)(a)(v), the medical advisory committee shall make recommendations about those issues to the hospital's quality committee established under subsection 3(1) of the *Excellent Care for All Act, 2010*.

The Quality Committee considers those recommendations under the Hospital Management Regulation, s. 7(8) when reporting to the board under section 3(3) of the ECFA. The Quality Committee is the independent statutory monitor of systemic quality on behalf of the board, and clinical AI quality at the systemic level is part of that monitoring function.

7.10 Reporting flows for clinical AI oversight

The reporting flow for clinical AI oversight could operate as follows:

- the AMC reports to the CEO on stream 2 deployment matters with MAC clinical input integrated through the chief of staff/MAC chair
- the CEO carries stream 2 recommendations to the Quality Committee at the materiality thresholds in the Sample Board AI Governance Policy
- the Quality Committee reports to the board on systemic AI quality under the Hospital Management Regulation, s. 7(8) and ECFA s. 3(3)
- the chief of staff/MAC chair reports MAC activities to the board, including on the practice of medicine under the Hospital Management Regulation, s. 7(5)

Critical incidents involving AI should follow any hospital internally adopted board notification requirements, such as the reporting obligations set out in the Sample Board AI Governance Policy, with regulatory disclosure obligations under the Hospital Management Regulation, s. 2(4) and the Cyber Security Regulation, s. 7.

7.11 Defending against cyber risk

Cyber attacks on hospitals have been a regularly recurring exposure for several years. Provincial response architecture has developed accordingly, including the Ontario Health Cyber Security Centre, sector-wide response teams, HIROC ransomware insurance coverage, and centralized incident response capability. That architecture exists because the threat is documented and active, and because the critical consequences of a hospital cyber incident, such as ambulance diversion, EMR downtime, deferred procedures, PHIPA breach exposure, and ultimately patient harm.

AI materially amplifies this pre-existing threat surface in two distinct ways. The first is risk arising from AI the hospital itself deploys. Each AI system increases the attack surface and creates new vendor and supply-chain dependencies. The second, and the focus of this section, is AI-enabled

cyber threats from adversaries operating against hospital infrastructure, regardless of the hospital's own AI choices. The second risk surface is stream-agnostic and arises whether or not the hospital deploys any AI. Threat actors are using AI capabilities against the hospital's existing infrastructure now.

The Ontario Health Cyber Security Centre's May 1, 2026 sector advisory describes the current operational reality: AI-enabled adversaries operate on compressed timelines that eliminate traditional warning signs.[97] Attacks are faster and more credible. Identity vulnerabilities are targeted first, and incident response plans calibrated to pre-AI attack patterns are no longer adequate. Defensive measures, including phishing-resistant multi-factor authentication on privileged, remote, and third-party access; elimination of shared, legacy, inactive, and test accounts; least-privilege access; network segmentation; rapid patching of internet-facing and high-risk systems; centralized monitoring of authentication, email, remote access, and cloud services; updated incident response plans accommodating AI-driven attack ambiguity; and clearly defined executive decision authority for rapid containment, should be prioritized. Several of these directives apply directly to staff use of personal phones and other personal devices to reach the EMR, email, and hospital meeting platforms.

Hospitals should consider having an AMC report on the hospital's personal-device access posture against these directives and on the hospital's response to the IPC's findings in *Re Otter.ai* (HR24-00691).

The scale of unsanctioned consumer AI access in public sector workplaces is documented. The Auditor General of Ontario found that, between April and August 2025, approximately 12,000 Ontario public service staff accessed roughly 400 AI websites, of which about 60% were rated unsafe or unsecured under Microsoft Defender's scoring, and that the Ministry had not implemented controls to prevent staff from uploading personal or sensitive information to those sites.[98] This is concrete evidence of the scale of shadow-AI exposure inside a public sector environment, and should inform the hospital's assessment of its own controls.

97 Ontario Health Threat Intelligence Advisory.

98 "Use of Artificial Intelligence in the Ontario Government: Independent Auditor's Report"

Article 8

The opportunities boards must enable

8.1 The governance obligation: risks and opportunities in equal measure

AI governance is not only a risk-management exercise. Under ONCA, directors must act in the best interests of the corporation, which for an Ontario hospital includes patients, staff, the Ministry of Health, Ontario Health, and the broader community.[99] A board that manages AI's risks, but fails to capture its benefits, would arguably fulfill only part of its fiduciary duty.

8.2 Stream 1: operational efficiency AI: The fiscal mandate answered

AI-optimized elective surgical scheduling has produced documented improvements in bed capacity utilization.[100] With more than 73,000 Ontarians waiting for surgery beyond clinically recommended times, even a 10% improvement in bed capacity utilization across Ontario's hospital system could help reduce the need for building new hospitals.[101]

8.3 Stream 2: clinical AI: the evidence-based strategy

Published evidence supports clinical AI deployment.[102] In the rural and remote care context, AI-enabled monitoring platforms extend specialist clinical capacity to patients, addressing the access inequity that is present in Ontario's healthcare environment. Published studies document significant reductions in hospitalizations through AI-enabled remote monitoring.[103]

AI platforms identifying high-risk patients before crisis presentation enable earlier, lower-cost intervention.[104]

99 ONCA, s. 43(1).

100 See, for example "Artificial Intelligence Length-of-Stay Forecasting and Pediatric Surgical Capacity", Barry J et al, JAMA Pediatrics, (March 2026).

101 "New report warns of longer wait-times, rushed care, and overcrowded Ontario hospitals as government cuts expected to cause over 10,000 job losses and shortfall of 4,080 beds", Canadian Union of Public Employees (CUPE). (January 27, 2026).

102 See, for example, footnotes 12, 13 and 14.

103 Medly.

104 See, for example, "Artificial intelligence in clinical risk prediction: promise, performance and the path forward?", Narasimhan P, BMJ Health & Care Informatics. (December 3, 2025).

8.4 Opportunity: workforce transformation and clinical capacity extension

A significant long-term institutional opportunity in Ontario is extending the effective clinical capacity of a workforce projected to lose 7,263 provincially-funded nurses by 2027-28.[105] AI documentation tools recover clinician time from administrative burden. AI-assisted diagnosis extends specialist scope beyond institutional walls. AI-enabled remote monitoring allows one specialist to serve patients across a large geographic catchment. These are workforce multipliers, not replacements.

The board's governance obligation extends beyond deploying AI tools to governing the workforce transformation that AI triggers. A strategic workforce transformation approach requires the board to oversee three questions:

- First, how is AI changing the skills the institution needs, and is management actively upskilling and reskilling staff to operate in an AI-augmented environment, rather than simply deploying tools on an unprepared workforce?
- Second, how are roles and job descriptions evolving, and is the institution prepared for the reality that AI changes what entry-level clinical and administrative positions require, which affects recruitment, onboarding, and the institution's ability to build a sustainable long-term talent pipeline?
- Third, where does AI support human judgment, and where does it risk eroding it? Boards should expect management to have an explicit human accountability framework that defines, for each AI deployment, which decisions must always involve meaningful human oversight, and which AI outputs may be acted upon without individual human review. AI that builds institutional capability is a governance success. AI that hollows out the clinical judgment of the workforce it was supposed to support is a governance failure.

8.5 Research AI: the academic hospital advantage

AI-generated intellectual property requires a robust governance model before research commences. By way of example of hospitals developing their own AI products, the University of Waterloo/Princess Margaret–UHN picosecond infrared laser mass spectrometry platform described in Schedule B illustrates a clinical research collaboration developed in an academic centre that, if successfully translated, could deliver real-time tumour identification during surgery and reduce the need for separate biopsy steps.

The governance imperative for academic hospitals is to build the framework that allows this kind of collaboration to translate from research into care under appropriate oversight, including REB review of research protocols, vendor and institutional accountability allocation between the academic partner and the hospital under PHIPA, AI and IP governance, appropriate guardrails to ensure privacy and health regulatory risks are being mitigated appropriately, and safe system design discipline is being implemented. In addition, intellectual property governance will be important where AI tools are used in research to develop potentially valuable assets,

105 Ontario Nurses Association. *"New report shows 7200 nurses could be cut in Ontario"*. October 28, 2025. ona.org

as current Canadian copyright and patent frameworks generally require human authors and human inventors, respectively, which may affect ownership, inventorship, and subsequent commercialization strategies.

Canada's national AI strategy, AI for All, reinforces this opportunity at the system level. The strategy recognizes that Canada's universal public health system is a distinctive national advantage, generating clinical and administrative data at a scale few jurisdictions can match, and commits funding — through platforms such as VITAL and the Health Sector Data Space — to mobilize patients' anonymized data responsibly to advance research, innovation, patient care, and system efficiency. That opportunity is expressly conditioned on safeguards that align with the obligations in this guide, including de-identification or anonymization of patient data, Canadian data residency, encryption, and access restricted to approved researchers.[106]

106 "Canada's National Artificial Intelligence Strategy: AI for All", Government of Canada, (June 4, 2026), Innovation, Science and Economic Development Canada, Pillar 4 (Unlocking the Power of Data) and the VITAL and Health Sector Data Space initiatives. On the de-identification and data-residency safeguards, see "Canada's AI strategy aims to make health data more accessible. How will it protect privacy?", Maia Tustonic, CBC News (June 24, 2026).

Article 9

AI procurement: the contractual framework

9.1 Procurement as an extension of governance

AI procurement is not simply a technology purchasing decision; it is an extension of the hospital's governance strategy. Decisions regarding which AI systems are adopted, how they are deployed, what data they can access, and the degree of autonomy, can directly impact patient safety, privacy, regulatory compliance and the hospital's broader strategic objectives.

Choosing an AI vendor does not transfer the hospital's legal accountability. The hospital remains the PHIPA health information custodian and, in most cases, will be responsible for ensuring devices are used and maintained in accordance with the manufacturer's authorized terms of use. AI vendor commercial interests may at times diverge or conflict with the board's governance obligations or strategic priorities. Accordingly, AI procurement must be viewed as a governance exercise that aligns technology adoption with the hospital's risk tolerance, objectives and priorities.

9.2 AI procurement must be use case driven

There is no single set of contractual terms that is appropriate for every AI procurement. The contractual protections, diligence requirements and governance measures that are appropriate for an AI-powered scheduling tool may be materially different from those required for an AI system that influences clinical decision-making or processes large volumes of personal health information. As such, contract terms must be calibrated to the specific use case, risk profile, and operating environment.

The hospital should ensure that an appropriate use case assessment is conducted before contractual terms are negotiated. At a minimum, the following questions should be considered.

- What is the use case? What is the AI tool intended to accomplish, and who may be affected by its use, including clinicians, staff or third parties?
- What data is involved? What types of data will be used by the AI tool (e.g., patient, employee, public data) and for what purposes, including inference, logging, model training or service improvement?
- What decisions or actions does it influence? Will the AI tool operate in an assistive, decision-support or autonomous capacity, and what decisions, actions or outcomes may it influence?
- What is the deployment model? Is the deployment a pilot or production implementation? What degree of customization, integration or development is required, and does the procurement involve software alone or a broader platform and related services?

9.3 Key contractual considerations for hospitals

Once the hospital has assessed the proposed AI use case and associated risk profile, it can evaluate the contractual protections required. Not every consideration discussed below will be relevant to every AI procurement. Rather, the depth and specificity of contractual protections should be proportionate to the nature of the use case, the data involved, the degree of autonomy, and the potential impact on patient safety, privacy, operations and regulatory compliance.

(a) Contractual terms

Based on the hospital's use case assessment and data governance review, consideration should be made with respect to every material AI vendor contract, proportionate to the risk with depth and specificity calibrated to the risk profile of the deployment. These considerations reflect regulatory expectations and best practice from applicable regulatory guidance and established healthcare AI contractual governance practice.

Contractual consideration	Risk profile relevance
Foundational definitions and governance alignment: Contractual definitions (e.g. what constitutes AI or AI commitments) should align with the hospital's governance framework. These definitions serve as the foundation for the balance of the contract	All material AI contracts.
Human oversight and override capability: For every AI system capable of influencing a clinical decision or otherwise posing a material risk to health and safety, a documented, tested, and operationally accessible mechanism must allow an individual, including a clinician where applicable, to override the AI output.	All Tier 2-4 AI risk clinical deployments at a minimum.
Health Canada medical device licence warranty: Where the AI system constitutes SaMD, the vendor should represent, warrant and covenant that all required Canadian regulatory authorizations have been obtained and will be maintained for the relevant time period, with an obligation for the vendor to notify the hospital of any material or significant changes.	All SaMD deployments.

Contractual consideration	Risk profile relevance
Data governance, PHIPA-compliance, and AI improvement rights: Patient data remains under the custody and control of the hospital. Vendor use of patient data and other data provided by the hospital is clearly defined and any use-rights limited, including restrictions on model training, fine-tuning, and product improvement purposes. Data location, subprocessors, retention and deletion obligations should also be addressed .	All AI systems processing personal health information.
AI system change and update management: Material changes to, or deprecation of, the AI system (including algorithm or model changes) are communicated to the hospital with sufficient lead time to assess clinical, operational, and regulatory impact before implementation.	All material AI contracts. Critical for Tier 2-4 AI risk where algorithm changes affect clinical outputs.
Performance commitments: Minimum data quality specifications should be defined and warranted by the vendor before deployment: training data must be representative of the deployment population; accuracy, completeness, consistency, and traceability must be documented; the vendor must disclose whether and how the system learns post-deployment; and the product's PCCP, if applicable, or the process for updating the model when performance degrades must be specified, with notification obligations to the hospital before any model update is deployed.[107] The vendor should provide sufficient documentation regarding the systems' intended use, limitations, testing methodology, and known performance constraints.	Tier 2-4 AI risk at a minimum.

107 See also, "International consensus guideline for trustworthy and deployable artificial intelligence in healthcare", BMJ, (2025), 388:e081554.

Contractual consideration	Risk profile relevance
Continuous post-deployment monitoring obligations: The vendor supports ongoing, not merely periodic, performance monitoring. The contract addresses notification of material performance issues, including where the AI system performs below predetermined accuracy thresholds, or produces unexpected outputs. Post-deployment monitoring should be mandatory, automated, and continuous and the hospital should understand the scope of the PCCP, if applicable.	All Tier 2-4 AI risk clinical deployments at a minimum.
Cybersecurity obligations: The vendor commits to security controls aligned with best practices, including prompt-injection and data-poisoning safeguards, AI bill of materials, vulnerability disclosure, identity controls, and audit rights. The vendor agrees to incident notification within timeframes that allow the hospital to meet its 72-hour Ministry CISO reporting obligations under the Cyber Security Regulation, and the hospital's internal board notification requirements, such as the 48-hour board notification set out in the Sample Board AI Governance Policy.	All Tier 2-4 AI risk deployments and any AI tool with EMR access or network integration at a minimum.
Explicit liability allocation: Liability is allocated to reflect each party's role and control over the AI system, including model design, updates, and operation. For agentic AI, liability for autonomous actions can be complicated and should be addressed in a manner that reflects each party's responsibilities and culpability.	All material AI contracts. Elevated importance for Tier 4 AI risk agentic AI deployments.
Exit rights, transition assistance and portability: Contractual terms should address termination, transition assistance, and portability of hospital data, outputs and hospital-developed AI assets. Hospital data should be provided in a usable format that aligns with the hospital's obligations under PHIPA. Where the hospital has invested in configuring, training or deploying AI systems, including agentic workflows or model customizations, the contract should address the hospital's ability to transition to alternative solutions to mitigate vendor lock-in.	All material AI contracts.

Contractual consideration	Risk profile relevance
Intellectual property and licensing considerations: Ownership of AI-inputs, AI-generated outputs, and model improvements is clearly set out. Generally, a hospital would expect to retain ownership of the data, content and other inputs that it provides to an AI system, together with sufficient rights to access, use and retain the resulting outputs for its clinical, operational, research and regulatory purposes. Where the hospital contributes to the development of model improvements, the parties should clearly address the ownership of, and licensing rights relating to, those improvements.	Academic hospitals deploying AI in research or discovery contexts.
Regulatory and standards alignment: For higher-risk use cases, the contract addresses vendor demonstration of alignment with relevant regulatory requirements and recognized standards.	Tier 3-4 AI risk and SaMD deployments where international best-practice alignment supports due diligence at a minimum.

(b) AI scribe procurement: Auditor General's findings and hospital diligence

The Auditor General's AI Report identified significant gaps in the provincial AI scribe procurement process conducted by Supply Ontario, the provincial government's supply chain agency, in consultation with Ontario Health and OntarioMD, an organization that supports family physicians and clinic staff to optimize technology use in their practices.[108] These findings are directly relevant to hospitals procuring AI scribe systems, whether through the provincial vendor of record arrangements or independently.

The report found that all 20 approved vendors in the Supply Ontario Vendor of Record arrangement demonstrated one or more of the following inaccuracies in their system-generated notes; hallucinations, incorrect information or incomplete information.[109] Forty-five percent of the AI scribe systems fabricated clinical information, 60% captured incorrect medications, and 85% missed key mental health details.[110] Security and accuracy criteria received disproportionately low scoring weight, and 11 of 20 vendors failed to submit third-party security reports, such as evidence of SOC 2 Type 2, HITRUST, or ISO 27001 certification.[111] Vendors were not required to provide bias testing evidence.[112] The process relied primarily on vendor self-attestation rather than independent verification.

108 "Use of Artificial Intelligence in the Ontario Government: Independent Auditor's Report".

109 Ibid at 23.

110 Ibid.

111 See "Tender-20123 – Artificial Intelligent Solutions – AI Scribe", Supply Ontario. (accessed May 2026). Supply Ontario, the provincial government's supply chain agency, has launched a Vendor of Record arrangement in which multiple AI Scribe vendors have been pre-qualified through a competitive request for bids process. The Supply Ontario website contains a full list of vendor profiles.

112 Ibid at 27.

A hospital procuring an AI scribe through the Supply Ontario Vendor of Record arrangement remains responsible for conducting independent vendor diligence sufficient to satisfy its legal obligations.[113] At a minimum, hospitals should ensure the following conditions have been satisfied and contractual protections are addressed:

- **Accuracy validation**. The vendor must provide evidence of accuracy testing using representative clinical scenarios and diverse patient populations. The contract must require IT controls enforcing clinician attestation before notes are finalized, define minimum accuracy thresholds with notification obligations, and must not rely solely on vendor self-attestation of accuracy.

- **Third-party security documentation**. The vendor must provide current SOC 2 Type 2 reports or equivalent certifications, such as HITRUST or ISO 27001, together with threat risk assessments and privacy impact assessments, as a precondition to contract execution. Annual provision of updated reports must be a contractual obligation.

- **Bias testing evidence**. The vendor must provide documented bias testing results demonstrating equitable performance across diverse patient populations, including varied accents, clinical contexts, and demographic groups. Descriptions of bias mitigation processes alone are insufficient.

- **Data residency verification**. The vendor should provide independent verification that all personal health information remains within Canada. While PHIPA does not categorically prohibit the storage of personal health information outside Canada, Canadian data residency materially simplifies compliance and reflects the prevailing sector expectation. Contractual terms must address the vendor's entire supply chain, including subprocessors and cloud infrastructure providers.

- **SOC report content review**. The contract must require vendor provision of SOC 2 Type 2 reports addressing AI-specific control objectives, and the hospital's evaluator must assess whether the report's control descriptions and any exceptions are responsive to the AI risks identified in the use-case assessment; not merely confirm that the report was provided.

- **Minimum passing scores and weighted evaluation criteria**. Where a hospital procures an AI scribe system outside the provincial vendor of record arrangement, the hospital's evaluation framework must assign minimum passing scores to security, privacy, accuracy and bias-control criteria, with materially higher weighting than the 11% (security) and 4% (accuracy) weights identified in the Supply Ontario request for bids reviewed by the Auditor General.

- **Live demonstration**. For any material clinical AI procurement, the hospital should require a live demonstration of the system performing the intended clinical task on representative use cases as part of its evaluation, consistent with the Auditor General's AI Report.[114] All AI scribes will not be designed to function effectively in all clinical environments.

113 "Tender-20123 – Artificial Intelligent Solutions – AI Scribe", Supply Ontario, (accessed May 2026). Supply Ontario, the provincial government's supply chain agency, has launched a Vendor of Record arrangement in which multiple AI Scribe vendors have been pre-qualified through a competitive request for bids process. The Supply Ontario website contains a full list of vendor profiles.

114 "Use of Artificial Intelligence in the Ontario Government: Independent Auditor's Report".

Article 10

The 12-month AI governance build roadmap

10.1 HAIRA AI governance maturity model

The Healthcare AI Governance Readiness Assessment (HAIRA) is a five-level model that identifies the maturity of healthcare institutions' AI governance practices from authors at the Centre for Computational Medicine and Clinical Artificial Intelligence, Department of Medicine, University of Chicago.[115] The five HAIRA levels are as follows:

(a) Level 1: initial/ad hoc

Hospitals at this HAIRA level have no formal AI governance and are operating with basic awareness and minimally structured processes. AI deployments are managed reactively, with no inventory, policy or reporting mechanisms. This is the starting point for many hospitals today.

(b) Level 2: developing

Hospitals at this level have a basic AI inventory and some policies in place. However, governance is informal and inconsistent. Hospitals have an awareness of regulatory obligations, but limited implementation.

(c) Level 3: defined

Hospitals at this level have adopted a formal AI governance policy and structured reporting is in place. In addition, the Quality Committee review process has been established and IRM integration has begun. HAIRA level 3 represents full compliance with the minimum governance architecture this guide describes.

(d) Level 4: managed

For hospitals at this HAIRA level, a comprehensive governance architecture is operational. Continuous monitoring is in place and the hospital is demonstrating full regulatory compliance. EDSTA compliance readiness has been achieved.

(e) Level 5: leading

Hospitals at this HAIRA level have AI governance embedded in the institutional culture. Such institutions have proactive risk management, full regulatory compliance, and are recognized externally as AI governance leaders.

115 "Advancing healthcare AI governance through a comprehensive maturity model based on systematic review", HAIRA, (February 11, 2026).

10.2 Implementation timeline

The following roadmap provides a sequenced illustrative implementation framework that could be applicable to all Ontario hospital types. The milestones are universal, and the timing and depth of implementation is flexible and proportionate to the institution's AI deployment profile, resources, and risk exposure. What is not illustrative is the obligation to have a roadmap, and to be executing against it.

The HAIRA level 3 governance baseline, completing an AI inventory, adopting a board policy, verifying SaMD authorizations as applicable, and establishing monitoring, should be achievable in the 12-month schedule set out below for every Ontario hospital.

Suggested timeline	Example milestone	Suggested owner
Months 1-3	Commission AI inventory of every AI tool deployed, classified by risk tier.	CIO/CTO
	Publish approved AI tools list and communicate employee AI Use Policy to all staff prohibiting use of non-approved AI tools for hospital work.	CIO/CTO
	Complete IPC AI scribe checklist for all deployed ambient AI documentation tools. Confirm CPSO patient consent protocol in place. Report compliance gaps to board.	Privacy officer and senior medical officer
	Board reviews and adopts AI governance policy (see the Sample Board AI Governance Policy). Legal counsel confirms legislative citations and materiality thresholds.	Board chair and legal counsel
Months 4-6	Update Quality Committee terms of reference to include AI governance mandate.[116]	Quality Committee chair and VP medical
	Verify Health Canada authorization for every AI tool obtained from a vendor qualifying as SaMD requiring a Health Canada licence. Document for Finance/Audit Committee.	Privacy officer, legal counsel and CIO/CTO

116 "Our Commitment to Using AI Safely, Responsibly, and Equitably", Stanford Health Care (accessed May 2026). Boards should consider the Stanford Health Care FURM (Fair, Useful, Reliable, Monitored) framework.

Suggested timeline	Example milestone	Suggested owner
Months 4-6	Launch institutional AI literacy program covering board, C-suite, and frontline IT/ clinical staff.	CEO and VP human resources, supported by the CIO/CTO
	Conduct HAIRA governance maturity assessment, discussed at Section 10.1. Report current level and targets to board.	CIO/CTO and AMC or designated stream committee
Months 7-9	Review top five material AI vendor contracts against key contractual considerations in the hospital's AI governance policy. Initiate re-negotiation for contracts with material gaps, if applicable. Plan for future contracts with AI vendors.	Legal counsel and CFO
Months 10-12	Establish semi-annual board AI governance reporting. Sample templates are set out in Appendices 1 and 2 to the Sample Board AI Governance Policy (Schedule D). First structured report to Quality Committee.	Senior medical officer and CIO/CTO
	Board skills matrix updated to include AI literacy and cybersecurity governance as sought competencies. Initiate multi-year board education curriculum. Confirm CEO accountability for C-suite AI fluency program.	Governance committee chair
	Annual AI performance and compliance report to board (Sample at Appendix 1 (annual report)). Conduct targeted compliance reviews for higher-risk AI deployments and conduct compliance review for ambient AI deployments.	CEO, senior medical officer and CFO

Suggested timeline	Example milestone	Suggested owner
By July 1, 2026	The Cyber Security Regulation obligations operationalized: designated cybersecurity primary contact and alternate identified to the Ministry CISO; baseline Cyber Security Maturity Assessment scoping initiated; 72-hour critical cybersecurity incident reporting workflow tested.	CIO/CTO and AMC
Within 12 months	AMC report to board on alignment with the CCCS *Top 10 AI Security Actions* (ITSAP.10.049) by pillar, with the CCCS *Cyber Security for Connected Medical Devices* primer (ITSAP.00.132) for SaMD deployments, and exception summary.	AMC and CIO/CTO

Article 11

Conclusion

AI is deployed in hospitals now. Hospital board's governance obligations are operative now. AI-enabled cyber threats are active against hospital infrastructure now.

Boards should approve an AI governance policy. The Sample Board AI Governance Policy establishes a structure for assessing, monitoring, and reporting on: (a) stream 1 operational AI deployment (b) stream 2 clinical AI deployment and (c) the management of AI risks, comprising both AI deployment risks across streams 1 and 2 and cyber risks, including AI-enabled cyber threats from external adversaries.

The principal components of the Sample Board AI Governance Policy are:

- An AI management committee constituted on the terms of reference described in Appendix 6 to the Sample Board AI Governance Policy (Schedule D).
- A current AI inventory maintained by the AMC and reviewed by the board annually, capturing every AI tool by risk tier, with Health Canada SaMD authorization verified for every qualifying deployment.
- A four-tier risk classification with named approval authority at each tier, ranging from CIO/CTO for Tier 1 AI risk operational AI to board approval for T4 AI risk agentic AI.
- A vendor contracting framework that addresses the considerations in Article 9.
- Template-driven structured reporting, including: an annual AI performance and compliance report (Appendix 1) and a semi-annual AI governance executive report (Appendix 2), with material incidents reported to the board within 48 hours (Appendix 4) and critical cybersecurity incidents reported to the Ministry CISO within 72 hours in accordance with the Cyber Security Regulation.
- Materiality thresholds defined covering patient harm, PHIPA breach, regulatory inquiry, financial exposure, media exposure, systemic shadow AI, and AI-enabled cyber threat indicators.
- A shadow AI detection and response framework (Appendix 5 to Sample Board AI Governance Policy) operationalising the IPC's findings in *Re Otter.ai* and the hospital's response to it.
- Director and C-suite AI fluency built through a multi-year curriculum tied to the institution's deployment profile, with the board skills matrix amended to identify AI literacy and cybersecurity governance as sought competencies.
- A board-level self-assessment checklist (Schedule E), including the Wyden Test, applied annually.

Proportionality governs the depth of implementation; not the existence of the obligation. A 1,200-bed academic health sciences centre and an 80-bed northern community hospital build the same components calibrated to different operating realities; engagement with Ontario Health shared-service models is the proportionate route at smaller institutions, and a deliberate complement at larger ones.

Boards that build this architecture are also positioned to capture the opportunity that Canada's national AI strategy, AI for All, identifies: a universal public health system whose anonymized data is a distinctive national asset for advancing research, innovation, patient care, and system efficiency, realized responsibly within the privacy and governance obligations this guide describes.

What this guide ultimately asks of an Ontario hospital board director is the question every director should be able to answer under oath, on the record, in June 2026: *I know what AI is deployed in my hospital. I know who is accountable for it. I know the governance framework that supervises it. I know what cybersecurity posture defends it. I receive structured reporting against that framework.*

The architecture that enables an Ontario hospital board director to answer those questions is the architecture this guide provides.

Schedule A:

Definitions

Capitalized terms used in this guide have the meanings set out below, listed in alphabetical order. Certain terms are defined directly in the guide where they first appear. Technical AI terms are intended for board directors without technical backgrounds. Legal and regulatory terms are shorthand references; full citations appear in the footnotes and body of this guide.

Agentic AI means AI that autonomously plans and executes multi-step tasks across multiple systems without human initiation of each step. Agentic AI represents the highest governance risk category and requires a board resolution before any deployment.

Algorithm means a set of instructions a computer follows to analyze data and produce an output. The algorithm is the engine inside an AI system.

Algorithm drift means the degradation of AI performance over time as clinical practice, patient populations, and care environments change without any change to the algorithm itself. Monitoring algorithm drift is a post-deployment AI governance obligation.

Ambient AI / AI scribe means a system that listens to a patient-clinician encounter and generates a draft clinical note for physician review and sign off. Use of ambient AI / AI scribes by hospitals is governed by the IPC's January 2026 AI Scribes Guidance and CPSO guidance, including patient consent requirements.

AMC means the proposed AI management committee established under Appendix 6 to the Sample Board AI Governance Policy (Schedule D). The AMC is a senior management body reporting to the CEO.

Artificial Intelligence (AI) means, as defined in EDSTA, a machine-based system that, for explicit or implicit objectives, infers from the input it receives to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments. This guide uses the EDSTA definition throughout.

Automation Bias means the tendency to over-rely on automated or AI outputs, reducing critical assessment and independent judgment. Automation bias should be addressed in mandatory training for clinical staff using AI decision support tools.

Board will have the meaning as set out in Section 1.1.

Business Judgment Rule (BJR) will have the meaning as set out in Section 3.2.

CCCS means the Canadian Centre for Cyber Security.

Chief of Staff / MAC Chair means the head of the medical advisory committee under PHA s. 35 and the Hospital Management Regulation, s. 7; an ex officio member of the board under hospital by-laws who reports to the board on the activities of the MAC and on the practice of medicine. This role is distinct from the senior medical officer.

CIO/CTO means the chief information officer or chief technology officer; the senior executive accountable to the CEO for stream 1 governance, AI inventory, the approved AI tools list, the shadow AI framework, and the cybersecurity operational alignment described in section 7.6. The CIO/CTO is the recommended designate for the Cyber Security Regulation cybersecurity primary contact.

CISO means the Ministry of Public and Business Service Delivery's chief information security officer.

CMA means cybersecurity maturity assessment.

CMIO means chief medical information officer, a senior employed clinical executive responsible for clinical informatics and digital-health/AI strategy, reporting to the CEO. The CMIO may serve on the AMC and attend the MAC ex officio non-voting.

ECFA means the *Excellent Care for All Act, 2010, S.O. 2010*, c. 14. Section 3 of the ECFA requires every hospital to have a Quality Committee chaired by a voting director and to monitor and report to the board on the quality of services provided in the hospital.

EDSTA means the *Enhancing Digital Security and Trust Act, 2024*. Schedule 1 to the *Strengthening Cyber Security and Building Trust in the Public Sector Act, 2024*, SO 2024, c 26. establishes AI accountability obligations for public sector entities. Enabling regulations are pending.

EMR means electronic medical record.

FURM means fair, useful, reliable, monitored. Stanford Health Care's publicly available pre-deployment AI evaluation framework.

Generative AI means a subset of AI that generates new content — including text, images, audio, video or software code — based on a user's request. Generative AI typically draws on large-scale training data to produce summarized or synthesized responses to queries. Examples include Microsoft Copilot Chat, ChatGPT, Claude and Gemini.

HAIRA means healthcare AI governance readiness assessment, a five-level maturity model (level 1: initial/ad hoc through level 5: leading) for evaluating institutional AI governance readiness.

Hallucination means when an AI system generates information that is factually incorrect but stated with apparent confidence. Hallucination is an inherent risk involved with current generative AI architectures.

HIROC means Healthcare Insurance Reciprocal of Canada.

Hospital will have the meaning as set out in section 1.1.

IPC means the Information and Privacy Commissioner of Ontario.

IRM means integrated risk management, the institutional framework for identifying, assessing, and monitoring organizational risks across all domains. AI should be added as risk sub-categories within each existing domain of every hospital's IRM framework, rather than as a separate IRM domain.

MAC means the hospital's medical advisory committee, constituted under PHA s. 35 and O. Reg. 965, s. 7.

MLMD means machine learning-enabled medical device. Any medical device that uses machine learning, in whole or in part, to achieve its intended medical purpose.

ONCA means the *Ontario Not-for-Profit Corporations Act, 2010*, SO 2010, c 15, the primary corporate governance statute for Ontario hospitals. The Act imposes fiduciary duties, duty of care, and the reasonable diligence reliance defence on directors and officers.

OPS means the Ontario public service.

PCCP means predetermined change control plan, a Health Canada-required document for MLMDs describing planned changes to the device and how they will be implemented and validated without requiring a new submission for each change.

PHA means the *Public Hospitals Act*, R.S.O. 1990, c. P.40, the provincial statute governing Ontario public hospitals.

PHIPA means the *Personal Health Information Protection Act, 2004*, SO 2004, c 3, Sch A, Ontario's health privacy statute. Every AI system accessing, processing, or generating personal health information engages PHIPA. The hospital, as health information custodian, is responsible for compliance with PHIPA.

Privacy impact assessment (PIA) means a risk management tool used to anticipate and address the potential privacy risks or impacts of a new or modified collection, use or disclosure of personal information, or of a proposed or existing information system, technology, program, process or other activity. Many institutions' policies require PIAs for projects involving personal information.

Quality Committee means the board committee established under ECFA s. 3(1) and O. Reg. 445/10, chaired by a voting director, that monitors and reports to the board on the quality of services provided in the hospital. For the purposes of this guide, the Quality Committee oversees stream 2 clinical AI deployments. May be an existing Quality or Quality of Care Committee with updated terms of reference.

Qmentum Global means Accreditation Canada's national hospital accreditation program, trusted by more than 9,300 health and social service locations across Canada.

REB means research ethics board.

RHPA means the *Regulated Health Professions Act, 1991*, SO 1991, c 18, Ontario's umbrella legislation governing regulated health professions.

SaMD means software as a medical device. Diagnostic AI, clinical decision support AI, and medical imaging AI are typically SaMD, software performing a medical function independently of hardware.

Shadow AI means any AI system used for hospital work not listed on the hospital's approved AI tools list. See Appendix 5 to Sample Board AI Governance Policy.

Senior Medical Officer means the senior employed clinical executive accountable for stream 2, typically the VP (Medical) or the chief medical information officer (CMIO), each reporting to the CEO. The senior medical officer role is distinct from the chief of staff/MAC chair.

Stream 1: operational efficiency AI means AI systems optimizing hospital operations without direct clinical decision-making: patient flow, surgical scheduling, bed allocation, staffing, supply chain, and back-office workflows. Stream 1 AI deployments tend to involve a lower governance risk than stream 2 deployments.

Stream 2: clinical AI means AI systems with direct clinical interfaces: deterioration prediction, diagnostic imaging analysis, clinical decision support, AI documentation tools, patient-facing AI, and research AI. Stream 2 AI deployments require evidence-based validation, Quality Committee oversight with senior medical officer recommendation, Health Canada SaMD authorization where applicable, and CPSO consent protocols.

Tri-Agency Framework means the joint policy of Canada's three federal research funding agencies (Canadian Institutes of Health Research, Natural Sciences and Engineering Research Council of Canada, and Social Sciences and Humanities Research Council) governing responsible conduct of research, including requirements regarding AI use in funded research.

Wyden Test means the test named for Senator Ron Wyden's questions at the U.S. Senate Finance Committee hearing on Change Healthcare. See Schedule E.

Schedule B:

AI examples and use cases

1. AI-enabled intraoperative cancer detection: ecosystem governance case study

The emergence of AI-enabled tools that operate in real time within clinical workflows represents a qualitative shift in the governance challenge facing boards. Unlike retrospective decision-support systems, which flag patterns in historical data for later clinician review, intraoperative AI compresses diagnosis, classification, and clinical action into a single step. The result is a category of technology that is simultaneously more clinically valuable and more difficult to govern than anything most hospital oversight structures were designed to address. This use of intraoperative AI is an example of the complexity of governing real-time AI-enabled tools and represents a near-term direction for clinical AI, a useful lens through which to illustrate the importance of the governance architecture set out in this guide.

Overview

Emerging clinical research illustrates how AI is moving from retrospective decision support into real-time, decision-adjacent care. University of Waterloo researchers, in collaboration with Princess Margaret/University Health Network (UHN) partners, are developing a picosecond infrared laser mass spectrometry diagnostic platform that generates molecular "fingerprints" of tumour tissue, with machine learning classifying high-priority central nervous system cancers in the operating theatre.[117] The platform is described as capable of identifying central nervous system cancer types in approximately 10 seconds to support intraoperative decision-making. By providing real-time classification of tumours during the surgery, the AI-enabled platform could eliminate separate biopsy and pathology steps and the need for additional surgeries to act upon the diagnostic findings.

Why this case study matters for boards

A complex AI deployment of this kind highlights five governance issues:

- **AI inside the clinical decision loop**. The AI output is generated and meant to be acted upon during the surgery, where decisions are immediate and irreversible. Governance must address risk at the point of care rather than retrospectively.
- **Accountability complexity**. The institution and the surgical team face dual exposure vis-à-vis failure to follow a validated AI output and inappropriate reliance on an AI output, with liability distributed across the surgical team, surgeon, anesthesiologist, pathologist, institution and vendor.
- **Standard of care**. If validated and adopted in health centres, such AI systems begin to influence what "reasonably prudent" Ontario hospital practice looks like, exactly the dynamic section 3.2 describes.

117 "Our Commitment to Using AI Safely, Responsibly, and Equitably", Stanford Health Care, (accessed May 2026). Boards should consider the Stanford Health Care FURM (Fair, Useful, Reliable, Monitored) framework.

- **Ecosystem risk**. The deployment is an interdependent system of hardware, software, clinical workflow, and institutional partners, none of which the hospital fully controls.
- **Data governance**. Molecular-level patient data is created at the point of care and engages PHIPA obligations on ownership, secondary use, and cross-institutional sharing that pre-existing data governance was not designed to address, creating a novel data governance question.

Governance controls

The patient and clinical benefits of an intraoperative diagnostic platform have governance, legal, and ethical risks, some of which are outlined below. Each pairing maps to a control already required by the architecture in Article 7 and the Sample Board AI Governance Policy.

Risk	Controls
Over-reliance on the AI output and compression of clinical judgment timeframe	• Human-in-the-loop oversight as described in section 4.1. • Mandatory clinician override capability documented in vendor contract per section 9.2.
Governance frameworks that are retrospective are ineffective for real-time clinical workflow	• Pre-approved, properly scoped use case approved by the hospital under the framework in section 7.1.
False positives or false negatives directly affecting irreversible surgical decisions	• Pre-deployment validation thresholds for sensitivity and specificity. • Continuous post-deployment monitoring against those thresholds. • A documented retraining or retirement pathway when thresholds are crossed.
Creation of novel personal health information and interpretability challenges	• PHIPA privacy impact assessment before deployment. • Express patient consent addressing molecular-level data, secondary use, and any cross-institutional data sharing. • Explainability obligations in vendor contract.

Risk	Controls
Drift, bias, and opacity of model outputs over time	• Continuous monitoring under layer 5 of the five-layer governance architecture set out in section 7.1(e). • Health Canada licensed SaMD terms and conditions are well understood, including any applicable PCCPs.
Undermining of independent clinical judgment in the time available to exercise it	• Independent local validation before clinical use. • Protocols defining how conflicts between AI output and clinician judgment are resolved, with the clinician as the final decision-maker.
Fragmented accountability across actors	• Contractual allocation of responsibility, indemnity, and incident-response duties across every party under the procurement framework in Article 9. • A single named operational and clinical owner inside the hospital under Appendix 3 to Sample Board AI Governance Policy.
Limited patient understanding of the AI's role in their care	• Disclosure and consent protocols meeting the CPSO standard. • Patient-facing materials describing what the AI does, what the clinician decides, and what data is created and retained.
Liability exposure from both misuse and non-use as the standard of care evolves	• Standing legal review of vendor contracts and clinical protocols. • Policy updates triggered by material model changes, regulatory developments, or jurisprudence under section 15 of the Sample Board AI Governance Policy.

Governance implication

This example illustrates that advanced clinical AI can concentrate clinical benefit, operational risk, accountability complexity, and patient trust within a single care pathway. A platform of this kind cannot be governed by any one instrument in this guide. Clinical AI of this nature requires the implementation of a full governance architecture, including the standard-of-care analysis under section 3.2, the duty-to-monitor obligations in section 3.3, the PHIPA framework in section 4.3;

Health Canada SaMD considerations under Article 6; the AMC and stage-gate discipline in section 7.1, the procurement and ecosystem-allocation framework in Article 9 and the materiality and incident-response thresholds, as illustrated in the Sample Board AI Governance Policy. The board cannot meet its obligations by approving any single element in isolation.

2. AI-enabled medical device examples

The following are examples of medical devices that use AI for illustrative purposes. Health Canada's Medical Devices Active Licence Listing does not separately identify AI/ML devices; so illustrative examples of medical devices available in the U.S. together with the FDA's short description on what the AI does in clinical practice, have been copied below from the U.S. FDA's AI-Enabled Medical Device List[118] as of May 10, 2026. The typical AI risk tier in accordance with section 7.3 of this guide that would likely be applicable to these devices also has been included for ease of reference. For the avoidance of doubt, the AI risk tier should not be confused with the class of a medical device, which is a separate classification that applies to all medical devices. FDA licence status is not relevant from a Canadian legal perspective, and all class II, III or IV medical devices imported or sold in Canada must be licensed by Health Canada.

a) Diagnostic imaging — radiology

For context, approximately 75% of FDA-authorized AI/ML devices are intended to be used for diagnostic imaging. The FDA's description of this category is AI reconstructs images from less raw data (faster scans, less radiation), measures specific anatomy, flags suspected abnormalities for the radiologist's attention, or routes urgent findings to the front of the worklist. These are typical Tier 3 AI risk deployments.

Manufacturer	Device	What it does	FDA
Siemens Healthineers	MAGNETOM Sola, Vida, Lumina	MRI scanners with AI-assisted image reconstruction	K252838
Hologic	Genius AI Detection 2.0	Marks suspicious lesions on screening mammograms	K243341
Hyperfine	Swoop Portable MR	Bedside low-field MRI with AI reconstruction	K253489

b) Stroke and neuroimaging triage

Time-critical AI that detects suspected stroke findings on CT and routes the patient and imaging directly to the on-call stroke team and interventional radiology, compressing scan-to-treatment time. The FDA's description of this category is an image processing device intended to aid in prioritization and triage of time-sensitive patient detection and diagnosis based on the analysis of medical images acquired from radiological signal acquisition systems. These are typical Tier 3 AI risk deployments.

118 Artificial Intelligence-Enabled Medical Devices, Food and Drug Administration (accessed May 2026).

Manufacturer	Device	What it does	FDA
iSchemaView	Rapid Aortic, Rapid OH, Rapid CTA 360	Detects large-vessel occlusion, perfusion deficits, aortic findings	K251987 K251533 K251151
Qure.ai	qER-CTA	CT angiography stroke triage	K251610

c) Cardiovascular AI

The AI analyzes ECGs to detect conditions a clinician may miss, simulates blood flow on a CT scan to avoid invasive cardiac catheterization, automates echocardiogram measurements, or analyzes ambulatory rhythm data at scale. The FDA's description of this category is this physiologic simulation software device is intended to aid in the identification of functionally significant cardiovascular disease. The following are typical Tier 2 AI risk to Tier 3 AI risk deployments.

Manufacturer	Device	What it does	FDA
Eko Health	Eko Foundation Analysis	Digital stethoscope with AI murmur and rhythm detection	K251494
Medicalgorithmics	DeepRhythmAI	Ambulatory rhythm AI analysis	K253141

d) Pathology AI

AI applied to digital whole-slide images to quantify biomarkers, classify tumour subtype, and generate prognostic risk scores that inform treatment selection. The FDA's description of this category is a device that provides prognostic risk estimates intended to assist physicians with prognostic risk-based decisions, but not intended to determine a clinical diagnosis. The following are typical Tier 3 AI risk deployments.

Manufacturer	Device	What it does	FDA
Roche Molecular Systems	Opulus Lymphoma Precision	AI-assisted lymphoma classification on digital pathology	K243863

e) Endoscopy and Polyp Detection

As part of this computer-aided detection during colonoscopy, the AI runs on the live video feed and highlights suspected polyps in real time, increasing detection rates compared with the endoscopist alone. The FDA's description of this category is a computer-assisted detection device used in conjunction with endoscopy for the detection of abnormal lesions in the gastrointestinal tract. The following are typical Tier 3 AI risk deployments.

Manufacturer	Device	What it does	FDA
Iterative Health	SKOUT	Real-time polyp detection during colonoscopy	K253664
Magentiq Eye	MAGENTIQ-COLO	Real-time polyp detection during colonoscopy	K252178

f) Ophthalmic AI

The AI analyzes retinal images and produces a diagnostic output without requiring an ophthalmologist to interpret. The FDA's description of this category is a prescription software device that incorporates an adaptive algorithm to evaluate ophthalmic images for diagnostic screening to identify retinal diseases or conditions. This is a typical Tier 3 AI risk deployment.

Manufacturer	Device	What it does	FDA
Eyenuk	EyeArt	Diabetic retinopathy detection	Multiple clearances
Carl Zeiss Meditec	CLARUS 700	Wide-field retinal imaging with AI analysis	K243878

g) Surgical and robotic AI

The AI segments anatomy from preoperative imaging to plan the procedure, registers patient anatomy to the robot intraoperatively, or guides instrument placement during surgery. The FDA's description of this category is an augmented reality device that provides visual guidance during procedures. The following are typical Tier 3 AI risk to Tier 4 AI risk deployments.

Manufacturer	Device	What it does	FDA
Mazor Robotics (Medtronic)	Mazor X / Mazor X Stealth	Robotic guidance for spine surgery (pedicle-screw placement)	K251316
Augmedics	xvision Spine	Augmented-reality spine surgery navigation	K251639
Procept Biorobotics	HYDROS	Robotic aquablation for benign prostatic hyperplasia	K251082

Schedule C:

Sample board resolution adopting the AI governance policy

Note: This is a starting framework; every provision requires review by legal counsel before adoption to ensure it is applicable to each hospital's specific circumstances. All legislative citations must be verified for currency. This document does not constitute legal advice and must not be adopted verbatim.

Resolution of the Board of Directors of

[**Hospital Name**] (the "**Hospital**")

Date: [**Date**]

[**Meeting: [Regular / Special] Meeting of the Board of Directors**]

Preamble

WHEREAS the board of directors of the hospital (the Board) wishes to approve an AI governance policy that will govern the deployment, monitoring and incident responses for AI at the hospital;

AND WHEREAS the Board was provided an overview of the following legislation and guidance obligations arising under the *Personal Health Information Protection Act, 2004*, SO 2004, c 3, Sch A; the *Enhancing Digital Security and Trust Act, 2024*, Schedule 1 to the *Strengthening Cyber Security and Building Trust in the Public Sector Act, 2024*, SO 2024, c 26; the Information and Privacy Commissioner of Ontario's guidance, *AI Scribes: Key Considerations for the Health Sector*; and the College of Physicians and Surgeons of Ontario's *Advice to the Profession: Using Artificial Intelligence in Clinical Practice*; as it relates to the deployment of an AI governance framework for the hospital;

Resolutions

BE IT RESOLVED THAT:

1. The Board hereby adopts the AI governance policy attached as Schedule [•] (the Policy) and directs management to implement the Policy in accordance with the implementation schedule agreed upon between the Board and management.
2. The Board confirms that operational accountability for AI deployment at the hospital rests with the CEO under O. Reg. 965, s. 3, advised by the AI management committee, and delegates to the Audit/Finance Committee the oversight of stream 1 AI deployments and to the Quality Committee, established under section 3(1) of the *Excellent Care for All Act, 2010*, the oversight of stream 2 AI deployments.
3. The Board acknowledges that the Policy is organized around two deployment streams and two risk categories:

a. Stream 1: operational efficiency AI, comprising tools that optimize patient flow, scheduling, bed allocation, staffing, supply chain, and administrative workflows without direct clinical decision-making, overseen by the Audit/Finance Committee;

b. Stream 2: clinical AI, comprising tools with direct clinical interfaces such as patient deterioration prediction, diagnostic imaging AI, clinical decision support, AI scribes, patient-facing AI, and research AI, overseen by the Quality Committee;

c. Risk 1: AI deployment risks arising from the hospital's own deployment of AI in streams 1 and 2; and

d. Risk 2: Cyber Risks comprising both AI-enabled cyber threats from external adversaries operating against the hospital regardless of the hospital's own AI choices and the cyber exposure created by AI systems the hospital itself deploys, with cybersecurity obligations integrated into the AI management committee's scope pursuant to the Cyber Security Regulation and reported through the committees overseeing both streams.

4. The Board directs management to ensure that the governance framework addresses both deployment streams and both risk categories with proportionate oversight, structured reporting, and defined accountability.

5. The Board directs the Governance Committee to update the Board skills matrix to identify AI literacy and cybersecurity governance as sought competencies, and to develop a multi-year director AI education curriculum, including programming to build and confirm C-suite AI fluency and institutional staff training protocols.

6. The Board directs management to add AI, where appropriate, as a named risk sub-category heading under the hospital's IRM framework.

CERTIFIED to be a true copy of a resolution passed [**at a meeting of the Board of Directors of the Hospital duly called and held on [Date], at which a quorum was present / by written resolution signed by all directors of the Hospital as of [Date]**].

_____

Board Chair

_____

Corporate Secretary

Schedule D:

Sample board AI governance policy

Note: This is a sample board-level governance framework for illustrative purposes. Policies are only effective if they can be realistically followed; so hospitals may wish to supplement and modify this governance policy to reflect their own institutional context, resources, and risk tolerance. Every provision requires review by the hospital's legal counsel, privacy officer, medical staff, and governance bodies before adoption. All legislative citations must be verified for currency. This document does not constitute legal advice. Management will need to develop operational policies, procedures, and protocols that give effect to this policy. This policy does not reflect all of the additional governance oversight that a hospital developing its own AI-enabled software and products would need to consider. Among other things, hospitals engaged in internal AI development will require more extensive policies addressing the software development lifecycle, model training and validation, technical risk management, and ongoing monitoring of internally developed systems.

1. Purpose and scope

This policy establishes the governance framework for artificial intelligence at [Hospital Name] (the Hospital) for the board of directors of the hospital (the Board) across clinical, operational, research, and educational functions. Management shall develop operational policies and procedures to give effect to this policy.

2. Guiding principles

All AI systems deployed by the hospital must be safe, effective, and ethical. For the purposes of this policy, AI deployments are organized into two streams, each with different governance imperatives, and the risks associated with AI are organized into two categories:

Stream 1: Stream 1 AI deployments are tools that optimize patient flow, such as scheduling, bed allocation, staffing, supply chain, and administrative workflows without direct clinical decision-making. These tools carry lower governance risk than stream 2 AI uses and should be implemented promptly. The hospital should assess stream 1 AI opportunities as part of its fiscal stewardship and, where appropriate, deploy them to help achieve the efficiencies Ontario's fiscal environment requires, consistent with the resource-allocation principles in section 3.5 of the guide.

Stream 2: Stream 2 AI deployments are tools with direct clinical interfaces such as patient deterioration prediction, diagnostic imaging AI, clinical decision support, AI scribes, patient-facing AI, and research AI. For all stream 2 deployments, patient safety is paramount, human clinical judgment is the final authority over all AI clinical outputs. These tools require evidence-based governance, phased deployment and clinical oversight.

Risk categories: The risks arising from AI fall into two categories: (a) *AI deployment risks*, comprising the patient safety, care quality, privacy, bias, automation bias, algorithm drift, accountability and vendor risks arising from the hospital's own deployment of AI in streams 1 and 2, and (b) *cyber risks*, comprising both AI-enabled cyber threats from external adversaries operating against the hospital regardless of the hospital's own AI choices, and the additional cyber exposure created by AI systems the hospital itself deploys. Boards should consider the risk of threat actors using AI against the hospital, whether the AI systems the hospital deploys are

protected from compromise, and whether the hospital's people and processes are positioned to use AI safely.

For all streams, AI use must be transparent, accountable, and privacy-protective. Equity in AI performance across demographic groups is a governance obligation, and all conflicts of interest must be disclosed and managed.

3. AI management committee

Operational accountability for AI deployment rests with the hospital's CEO under O. Reg. 965 to the *Public Hospitals Act*, s. 3. The hospital's CEO is advised by the AI management committee (AMC). The terms of reference for the AMC are attached as Appendix 6 to the Sample Board AI Governance Policy. The Board delegates AI governance oversight at the Board level to existing stream committees: the Audit/Finance Committee for stream 1 and the Quality Committee, established under the *Excellent Care for All Act*, s. 3(1), for stream 2.

The CEO, or designate, is the senior employee supporting both committees. Both committees receive structured semi-annual reporting and conduct a joint annual cross-stream review covering the AI inventory, this policy, and regulatory compliance. The Board directs the CEO to establish the AMC as a management-level body reporting to the CEO. The AMC executes the hospital's AI governance program operationally and reports through the CEO to the relevant stream committees. AMC's composition shall be comprised of, at a minimum, the:

- CIO/CTO (chair)
- Senior medical officer (VPM or CMIO)
- Chief privacy officer
- Legal counsel
- Ethics advisor
- Risk management designate
- CFO
- VP human resources
- Chief of staff or designate from the medical advisory committee (MAC)
- Frontline nurse or allied health professional, and
- AI or informatics specialist.

The designated primary contact for the purposes of the Cyber Security Regulation must be an AMC member.

4. AI risk classification, approval authority

The hospital applies a four-tier AI risk classification, with governance rigour increasing from Tier 1 (lowest risk) to Tier 4 (highest risk). The tiers are distinguished by the AI system's proximity to clinical decision-making and its degree of autonomy, and are defined as follows:

- **Tier 1 AI risk – operational efficiency AI**: AI that supports hospital operations without direct clinical decision-making, such as patient flow optimization, surgical scheduling, bed allocation, staff scheduling, supply chain, billing, and administrative workflows.

- **Tier 2 AI risk – clinical decision support AI**: AI that informs a clinical decision without generating autonomous clinical output, such as patient deterioration monitoring, readmission-risk prediction, and diagnostic flagging presented for clinician review.

- **Tier 3 AI risk – autonomous diagnostic AI / SaMD**: AI that generates autonomous diagnostic outputs or clinical recommendations, such as medical imaging AI and regulated AI-enabled medical devices (software as a medical device).

- **Tier 4 AI risk – agentic AI / high-risk SaMD**: AI that autonomously executes multi-step workflows across clinical systems without per-step human authorization, together with the highest-risk software as a medical device.

All AI systems must be classified by AI risk tier before deployment, applying the four-tier AI risk classification framework set out in this policy. Approval authority:

- Tier 1 AI risk: CIO/CTO approval

- Tier 2 AI risk: AMC recommendation, with CIO/CTO approval (CFO concurrence) for stream 1 and senior medical officer endorsement and CEO approval for stream 2

- Tier 3 AI risk: senior medical officer endorsement and CEO approval, with privacy officer and legal counsel review, and

- Tier 4 AI risk: board of Directors resolution required before any deployment.

No AI system may be deployed without approval at the required level. AI risk tier classification and approval must be documented in the AI inventory before deployment proceeds. For Tier 3 and Tier 4 AI risk deployments, independent validation of performance and safety — independent of both the vendor and hospital management — shall be obtained before approval and at defined re-validation points. Approval of a Tier 3 or Tier 4 AI risk deployment shall be supported by a short (one-page) deployment safety case that identifies the deployment's principal risks and states, for each, the mitigation relied upon. A material change to a deployed tool's autonomy, intended use, or underlying model may require re-classification and re-approval at the revised tier.

5. AI Inventory

The CIO/CTO shall maintain a current AI inventory listing every AI system deployed by the hospital, its AI risk tier, stream classification, hospital approval status, Health Canada licence status (for software as a medical device (SaMD)), and for licensed SaMD, the scope of the product's terms and conditions, including whether the software has a PCCP, the expected update frequency, and post-deployment monitoring status.

The AMC reviews the inventory semi-annually; the board reviews it annually. The inventory captures all deployed AI, including tools deployed without proper approval, which are flagged for immediate remediation. The AI management committee shall be responsible for approving the AI inventory template to ensure standardized reporting (see Appendix 3).

6. Employee AI use and shadow AI governance

No employee, physician, contractor, or volunteer may use any AI system for hospital work unless that system is on the hospital's approved AI tools list. Use of any unlisted AI system is prohibited, including personal consumer AI accounts and departmentally procured tools not reviewed through the governance process. The CIO/CTO shall publish and maintain the approved list, updated within 30 days of any change. Management shall develop an operational employee AI use policy, a voluntary self-reporting pathway, and an educational-first response framework. Any unauthorized AI use involving patient health information triggers immediate *Personal Health Information Protection Act, 2004*, SO 2004, c 3 Sched A (PHIPA) breach assessment. See Appendix 5 to Sample Board AI Governance Policy.

7. Prohibited uses

The following uses are prohibited under Ontario law regardless of Board authorization: AI systems that communicate a diagnosis directly to a patient without regulated professional review, contrary to RHPA s. 27(2)1.

The hospital also prohibits the following under this policy, whether or not otherwise unlawful in Ontario: (a) any AI system designed to manipulate patient or staff behaviour through deceptive techniques causing or likely to cause harm (b) any AI system exploiting vulnerabilities arising from age, disability, or specific social or economic situation to materially distort the behaviour of a patient or staff member in a manner causing or likely to cause harm (c) any AI system used by the hospital to categorise or score patients or staff on the basis of protected characteristics under the *Human Rights Code*, RSO 1990, c H.19 where doing so produces detrimental or unfavourable treatment unconnected to the purpose for which the data was generated, and (d) any AI system using real-time remote biometric identification of patients or members of the public in publicly accessible areas of the hospital, except as required for explicitly authorised patient identification or safety purposes documented in policy.

The following uses are prohibited without specific Board approval and legal counsel review: (a) biometric categorization systems using sensitive characteristics (b) AI systems inferring emotional states without express consent, and (c) agentic AI systems operating autonomously across multiple clinical systems without per-step human authorization.

8. Regulatory compliance

Stream 1: operational AI. Every AI system processing personal health information, regardless of whether its function is clinical or operational, should have a PHIPA privacy impact assessment before deployment. Enhancing Digital Security and Trust Act, 2024, SO 2024, c 26, Sched 1 (EDSTA) accountability obligations apply to all AI systems deployed by the hospital; and management shall build compliance readiness in advance of enabling regulations. No Tier 1 AI risk operational tool may be deployed without AI tier classification, approval, and listing on the approved AI tools list.

Stream 2: clinical AI. Software should be reviewed to confirm whether it is considered SaMD requiring a licence issued by Health Canada and, if so, licence status should be confirmed and documented. Patient consent protocols for ambient AI documentation deployments must meet the College of Physicians and Surgeons of Ontario requirements: consent must be express, informed, and recorded in the health record. Management shall develop policies governing all AI

documentation tools that are compliant with Information and Privacy Commissioner of Ontario requirements and guidance. No Tier 2, 3 or 4 AI risk clinical tools may be deployed without AI tier classification, approval in accordance with the risk classification and approval authority, and listing on the approved AI tools list.

9. Accountability structure

The Board is accountable for: adopting and maintaining this policy; approving Tier 4 AI risk deployments by resolution; receiving annual AI performance and compliance reports; director AI literacy through a structured education curriculum; and holding the CEO accountable for the hospital's AI governance framework consistent with the CEO's statutory obligations.

The CEO is accountable for implementing this policy and ensuring C-suite AI fluency across all executives with AI governance responsibilities. The AMC reports to the CEO and advises on AI deployment matters, with onward flow through the CEO to the Audit/Finance Committee for stream 1 and the Quality Committee for stream 2.

The CIO/CTO is accountable to the CEO for: stream 1 governance; AI inventory; approved AI tools list; shadow AI framework; and stream 1 monitoring and reporting. The CIO/CTO, together with the senior medical officer, is accountable to the CEO for: stream 2 clinical AI governance; clinical oversight; and stream 2 clinical performance monitoring and reporting.

The privacy officer is accountable for PHIPA compliance across both streams, privacy impact assessments, and breach reporting.

The MAC retains its statutory function under O. Reg. 965 to the Public Hospitals Act, s. 7:

- individual physician privilege recommendations flow MAC-to-Board under *Public Hospitals Act* s. 36 and s. 37
- systemic or recurring quality of care recommendations flow MAC-to-Quality Committee under O. Reg. 965, s. 7(7)
- the MAC's clinical input on proposed AI deployments is sought through the CEO's management process and integrated into the AMC's deliberations.

Management shall assign operational accountability within these governance structures through operational policies.

10. Staff training and AI fluency

The hospital shall maintain mandatory, role-specific AI training for all staff with access to hospital systems, with completion tracked and reported annually to the Board. Completion of role-appropriate training is a precondition to access to any hospital AI system. Refresher training shall be provided at minimum every 24 months and following any material change in deployed AI systems or applicable law. The hospital's training is not voluntary.

The Board shall ensure a structured director AI education curriculum. The CEO shall ensure C-suite AI fluency appropriate to each executive's governance responsibilities. Management shall develop

operational training protocols addressing: role-specific training calibrated to AI risk tier; mandatory automation bias awareness training for clinical staff; training documentation as a precondition to system access for Tier 2-4 AI risk deployments; mandatory retraining on material system updates; and competency review linked to post-deployment monitoring findings.

11. Vendor governance

All material AI vendor contracts shall address, at a minimum, the contractual considerations in this policy, calibrated to risk profile. No material AI contract may be executed without legal counsel review. Validation of Tier 3 and Tier 4 AI risk clinical AI shall not rely solely on vendor self-attestation; independent verification is required. Where AI is procured within a long-duration or multi-vendor transformation structure, including operate-to-transform arrangements in which private-sector capital modernizes hospital systems over multi-year terms, the use-case assessment shall specifically address exit rights, data and AI-capability repatriation, offshore processing, and vendor lock-in. Human oversight and override capability is mandatory for every AI system with direct clinical application. Procurement shall begin with a use-case assessment establishing: data sensitivity; clinical or operational impact; system autonomy level; and patient safety, privacy, and regulatory risks. Management shall develop vendor management protocols and contract templates consistent with this framework.

12. Monitoring and reporting

The hospital shall maintain a current AI inventory covering all deployed AI tools, reviewed by the AMC semi-annually and by the Board annually. Stream 1 monitoring tracks: efficiency metrics against objectives; approved tools list currency; shadow AI incidents; and EDSTA compliance readiness.

Stream 1 monitoring shall also track positive adoption metrics for approved AI tools, including unique active users by role, departmental adoption rates, and the ratio of approved-tool usage to known shadow AI incidents. The hospital shall set internal adoption targets for each material approved tool and report variance to the Audit/Finance Committee semi-annually. An approved tools list with negligible adoption is operationally equivalent to no approved tools list.

Stream 2 monitoring shall be continuous for all Tier 2-4 AI risk tools, tracking: clinical performance against benchmarks; demographic subgroup performance; algorithm drift; clinician override rates; and adverse events and near-misses. Stream-agnostic cyber monitoring tracks the hospital's posture against the CCCS *Top 10 AI Security Actions* (ITSAP.10.049), the CCCS *Cyber Security for Connected Medical Devices* primer (ITSAP.00.132) for SaMD deployments, the Cyber Security Regulation compliance status (designated primary contact, Cyber Security Maturity Assessment, and 72-hour critical incident reporting), and any provincial health-sector cyber advisories acted on since the last report.

The CIO/CTO reports through the CEO to the Audit/Finance Committee semi-annually on stream 1 and on stream-agnostic cyber posture. The senior medical officer reports through the CEO to the Quality Committee semi-annually on stream 2. Both use templates in Appendix 1 and Appendix 2 to this policy. An annual AI performance and compliance report covering both streams and the cyber posture is to be delivered to the Board each fiscal year.

13. Materiality thresholds and incident response

An AI-related event requires Board notification within 48 hours of management confirming that it:

(a) resulted in patient harm or near-miss attributable to AI

(b) constituted a PHIPA breach involving more than [**THRESHOLD**] records

(c) triggered a regulatory inquiry or enforcement action

(d) created financial exposure exceeding [**THRESHOLD**]

(e) is reasonably likely to attract media coverage

(f) involves shadow AI use at systemic scale, or

(g) involves an AI-enabled cyber incident or credible threat indicator with potential patient safety, operational continuity, or PHIPA exposure, including incidents arising from threat-actor AI capabilities operating against hospital infrastructure independent of hospital AI deployment.

Management shall develop an AI incident response protocol operationalizing this obligation, using Appendix 4 (material incident report). The Board's role in incident response is oversight, not management: the CEO manages the incident; the Board receives structured notification. Hospital incident response policy addressing privacy breach or incident response more generally could be cross-referenced, as applicable.

14. Conflict of interest

AI vendor relationships are a declared conflict of interest under this policy. All Board members, MAC members, and senior leadership must disclose AI vendor relationships in the annual conflict of interest disclosure process under s. 41 of the *Ontario Not-for-Profit Corporations Act, 2010*, SO 2010, c 15. Disclosure triggers recusal from all governance decisions involving that vendor. Management shall develop an operational conflict of interest screening protocol for AI procurement decisions. Hospital procurement policy addressing conflict of interest more generally could be cross-referenced, as applicable.

15. Policy review

This policy shall be reviewed no less than every two years, or earlier following: a material AI-related incident; a significant regulatory development, including EDSTA-enabling regulations; a material change in the hospital's AI deployment profile; or a significant development in Ontario AI healthcare liability jurisprudence. The CEO shall recommend amendments to the Board. The Board approves all amendments to this policy.

Appendix 1 to sample board AI governance policy:

Annual AI performance and compliance report template

This is a sample board-level reporting template for illustrative purposes. Hospitals should seek legal advice prior to adopting this reporting template.

Instructions: This template is for the annual AI performance and compliance report delivered to the board each fiscal year. Complete all sections. Attach supporting documentation where indicated. This report is a board governance document. As such, it should be minuted and retained. Report each metric against the prior year to show trend, and summarize at the front the decisions the board must make this year.

Section 1: AI inventory summary

Category	Details
Total AI tools deployed (as of report date)	
Tier 1 AI risk tools (operational efficiency AI)	
Tier 2 AI risk tools (clinical decision support)	
Tier 3 AI risk tools (autonomous diagnostic AI / SaMD)	
Tier 4 AI risk tools (agentic AI / high-risk SaMD)	
New deployments in the past 12 months	
Decommissioned or paused tools in the past 12 months	

Section 2: Regulatory compliance status

Compliance item	Status and notes
Health Canada SaMD authorization: all qualifying tools verified?	
IPC AI scribe checklist completed for all ambient AI documentation tools?	

Compliance item	Status and notes
CPSO patient consent protocols in place for all physician AI scribe deployments?	
PHIPA privacy impact assessments completed for all AI processing PHI?	
EDSTA compliance readiness plan in place?	

Section 3, stream 1: operational AI performance summary

3.1 Approved AI tools list

Current as of: [**date**]	Last updated: [**date**]
Tools listed: [**number**]	Changes since last report: [**number of additions / decommissions**]

3.2 Shadow AI incidents

Level 1 [**number**]	Level 2 (patient data) [**number**]	Level 3 (repeat) [**number**]	Level 4 (systemic) [**number**]
PHIPA breaches: [**number**]			
Remediation: [**summary**]			

3.3 Operational performance

For each material stream 1 tool: Tool | Objective | Metric | Result vs. Target | Variance

3.4 EDSTA readiness

Status: [**current / in progress / not commenced**]
Timeline: [**date**]
Gaps: [**summary**]

3.5 Material procurements contracts

Executed, renewed, or under negotiation: [summary]

Section 4, stream 2: clinical AI performance summary

Performance item	Details
AI tools with active continuous post-deployment performance monitoring	
Tools flagged for performance degradation or algorithm drift in past 12 months	
Demographic subgroup monitoring in place (Y/N per tool)	
AI-related adverse events or near-misses reported	
OHIP billing anomalies detected and addressed	
Independent validation completed for Tier 3-4 AI risk tools (Y/N per tool; date)	

Section 5: vendor contract compliance

Contract item	Status
Number of material AI vendor contracts reviewed against key contractual considerations	
Contracts requiring renegotiation identified	
Human oversight and override capability confirmed in all Tier 2-4 AI risk contracts	
AI system change and update management clauses confirmed in all material contracts	

Section 6: governance process compliance

Governance item	Status
Board AI education provided in past 12 months (date, topic, participants)	
AI governance competency identified in board skills matrix	
C-suite AI fluency program implemented and current	
Staff AI training protocols in place and documented	
AI named as a risk category in IRM framework	
AMC meetings held (number, key decisions)	
HAIRA maturity level (current and target)	
Conflict of interest disclosures received and any recusals	

Section 7: Key risks and opportunities

Instructions: Provide 2-3 sentences for each of the following. Include the management response taken or governance direction given.

- Summarize the top three AI governance risks identified in the past 12 months and management's response.
- Summarize the top three AI-enabled opportunities identified and the board's strategic direction for deploying AI safely, effectively and ethically to drive system-level clinical, efficiency and economic benefits.

Section 8: Recommendations to the board

List specific recommendations requiring board direction, approval, or awareness.

Section 9: Attestation

Prepared by:	Title:
Date:	Reviewed by Legal Counsel:

Appendix 2 to sample board AI governance policy:

Semi-annual AI governance executive report template

This is a sample board-level reporting template for illustrative purposes. Hospitals should seek legal advice prior to adopting this reporting template.

Instructions: This template is for the semi-annual AI governance executive report delivered to the board (or designated committee) every six months. It focuses on material developments, emerging risks, and performance trends since the last reporting period. This semi-annual report supplements, but does not replace, the annual AI performance and compliance report (Appendix 1 (annual report)). Where both templates address the same domain, Appendix 1 (annual report) provides the annual baseline inventory and compliance status; this template provides the interim update.

Reporting Period	FROM:	TO:
Prepared by:		
Date delivered to Board/Committee:		

1. **Board oversight**. Summarize the specific decisions the board must make this cycle, with the supporting detail provided under the relevant domain below. Also identify the items requiring only the board's awareness.
2. **Overall AI risk posture**. State the overall AI risk posture this cycle — improving, stable, or deteriorating — with a one- to two-sentence explanation of the direction of travel across the AI portfolio, ahead of the domain-level reporting below.
3. **Board reporting pack: five reporting domains**

This report covers five domains. Each domain should be reported concisely. The board should oversee value, technology exposure, control health, workforce accountability, and material incidents; not the day-to-day operation of AI models. Each domain should be reported against the prior period so the board sees trends; not only a point-in-time snapshot, and names an accountable executive: Domain 1, the CIO/CTO; Domain 2, the senior medical officer and privacy officer; Domain 3, the CIO/CTO; Domain 4, the CEO; Domain 5, the CFO.

Domain 1: AI deployment footprint. Active use cases by risk tier, domain, and lifecycle stage; new deployments and decommissions since last report; pilot vs. production distribution.

Domain 2: Trust, safety and privacy. Material incidents, near-misses, and overrides since last report; PHIPA breaches involving AI; shadow AI incidents by level; patient complaints involving AI.

Domain 3: Technology and third-party exposure. Critical vendor dependencies; material model changes, outages, or performance notifications; security posture and concentration risk; SaMD authorization status for all qualifying tools; AI-enabled cyber threat posture: alignment with CCCS's *Top 10 Artificial Intelligence Security Actions* (ITSAP.10.049), exception summary, and any provincial health-sector cyber advisories acted on since the last report.

Domain 4: Control and assurance health. Open approval conditions; overdue reviews; control failures; assurance findings; stream 1 EDSTA readiness status; stream 2 Quality Committee oversight status, including Tier 3 CEO approvals and any senior medical officer endorsement escalations to the Quality Committee.

Domain 5: Value, investment and workforce readiness. Realized operational benefits for stream 1 tools against deployment objectives; major AI-related expenditure; staff training completion rates; human accountability gaps identified.

4. Immediate escalation triggers

Board notification within 48 hours (not deferred to next semi-annual report) is required when: (1) a patient safety, privacy, cyber, or compliance event creates material exposure (2) a material vendor change, outage, or concentration risk alters the institution's risk profile (3) a key control fails, local validation is absent, or assurance identifies a material weakness (4) a major capital decision, strategic pivot, or workforce impact falls outside approved guardrails.

5. Material developments since last report

List any new AI deployments, decommissions, significant algorithm updates, vendor changes, or regulatory developments since the last reporting period.

6. Incidents and near-misses

Incident category	Number / details	Prior period
AI-related adverse events		
AI-related near-misses		
PHIPA breaches involving AI systems		
Regulatory inquiries involving AI		
Vendor notification of performance issues		

7. Tier 3 approvals and exceptions log

Log every Tier 3 deployment approved by the CEO; each instance in which the senior medical officer declined to endorse a Tier 3 deployment and the matter was escalated to the Quality Committee; and any approved exception to a board guardrail. For each entry, record the date, the deployment, the approver, and the rationale; so the board holds a contemporaneous record of its oversight.

8. Tier 4 approvals

For each Tier 4 system in operation, confirm the authorizing board resolution and its date, and report compliance status against every condition the board attached to that resolution. List any system classified, or proposed for classification, as Tier 4 that is under evaluation and will require a board resolution before deployment, with the anticipated approval cycle. Any Tier 4 system found in operation without a supporting board resolution is a control failure requiring immediate board notification under the escalation triggers above.

9. **Performance monitoring update**

Summarize performance monitoring results for all active Tier 2-4 AI risk tools. Flag any tools approaching or below performance thresholds. Report on demographic subgroup monitoring findings.

10. **Regulatory and compliance update**

Report on any new regulatory guidance, legislation, or enforcement actions affecting the hospital's AI governance obligations since the last reporting period.

11. **Emerging risks and opportunities**

Identify any new AI tools under evaluation, new vendor relationships, or operational changes that create new or elevated governance risk. Identify emerging AI-enabled opportunities and the governance steps required to enable those opportunities safely, effectively and ethically.

Appendix 3 to sample board AI governance policy:

AI inventory template

This is a sample AI inventory template for illustrative purposes. Hospitals should seek legal advice prior to adopting this reporting template.

Every Ontario hospital should maintain a current AI inventory as the operational record system underpinning all AI governance reporting. The inventory is an accountability instrument; not a catalogue. An entry is not complete unless all seven minimum fields are populated for each deployed or pilot AI system. The AMC is accountable for inventory integrity and must confirm completeness before every semi-annual board report.

Note: Complete the following for each AI tool.

1. **Use Case ID and lifecycle status**. Assign a unique identifier. Record current lifecycle stage: Concept | Pilot | Production | Under Review | Retired. Date of each stage transition must be logged.
2. **AI risk tier and supporting assessment**. AI risk tier classification (1-4) with reference to the risk assessment document that supports it. AI risk tier classification must be reviewed whenever the use case scope or the AI system materially changes.
3. **Operational and clinical owner**. Named individual accountable for ongoing performance (operational owner, typically CIO/CTO for stream 1, senior medical officer or designated physician lead for stream 2). Both owners must be named for stream 2 deployments. Owner name and title; not position description only.
4. **Vendor and contract reference**. Vendor name, contract number, and contract expiry date. Health Canada licence status: Licensed SaMD (licence number) | Not Required (confirm basis) | Pending | Not Confirmed. FDA clearance or authorization must not be recorded as it is not relevant or equivalent to Health Canada licensure.
5. **Approval conditions and compliance status**. List each condition attached to deployment approval (e.g., mandatory clinician review before action, quarterly demographic subgroup reporting, limitations and updates identified in the product's terms and conditions are diarized). For each condition: Met | In Progress | Overdue. Overdue conditions must be flagged in the next board report.
6. **Monitoring summary**. Date of last performance review. Summary of findings (Performing to Standard | Performance Concern Identified | Under Review). Number of clinician overrides since last review. Number of AI-related incidents or near-misses since last review. Demographic subgroup performance status (Current | Overdue | Not Applicable).
7. **Next review date and change-control linkage**. Scheduled next review date. Any pending change-control items (vendor model update notified, scope change under assessment, contract renewal due). Change-control items must trigger a reassessment before the change is implemented.

AI risk register entries: technology and third-party exposure domain

In addition to the per-tool inventory entries, the hospital's AI risk register under the Technology and Third-Party Exposure domain (Appendix 2 to Sample Board AI Governance Policy) shall include the following enterprise-level AI risk entry, separately from any deployment-specific risks:

- **AI-enabled cyber threat amplification**. Material increase in the speed, scale, and credibility of cyber-attack vectors targeting hospital infrastructure (phishing, impersonation, vulnerability weaponization, reconnaissance acceleration), arising from threat-actor deployment of AI capabilities and operating independently of the hospital's own AI deployment choices. Mitigations: alignment with the CCCS's *Top 10 Artificial Intelligence Security Actions* (ITSAP.10.049) operational framework; participation in Ontario Health Cyber Security Centre response architecture; AMC oversight of AI-enabled threat posture; materiality threshold for board notification. Owner: CIO/CTO, supported by AMC. Reporting: semi-annual under Appendix 2.

- **Board and executive dashboard view**. The board does not receive the full inventory; it receives a structured summary derived from it. The semi-annual AI governance report (Appendix 2) must present the following dashboard fields:

  1. Deployment footprint: total active use cases by risk tier (Tier 1-4) and stream (stream 1 or 2), pilot vs. production split
  2. Trust, safety and privacy: number of material incidents, near-misses, and clinician overrides since last report
  3. Technology and third-party exposure: vendor concentration, material contract renewals due, any vendor-notified model changes
  4. Control and assurance health: open approval conditions, overdue reviews, failed controls
  5. Training and workforce readiness: training completion rates for Tier 2-4 AI risk deployments, material training gaps. The immediate escalation triggers set out in Appendix 2, Section 2 apply.

Integrity standard

The AMC must confirm to the relevant stream committee, before each semi-annual board report, that: (a) every AI tool currently deployed or in pilot is listed (b) all seven fields are populated for each entry (c) ownership fields name individuals; not positions (d) all overdue conditions and reviews are flagged, and (e) the inventory has been reviewed by legal counsel and the chief privacy officer within the past six months. An inventory that has not been confirmed under this standard does not support the board's oversight obligation. The board's governance question is not only what the inventory shows; it is how the inventory was compiled and verified.

Appendix 4 to sample board AI governance policy:

Material incident 48-hour report template

This is a sample board-level reporting template for illustrative purposes. Hospitals should seek legal advice prior to adopting this reporting template.

Instructions: This template is for the board notification required within 48 hours of management confirming a material AI-related incident. It is a preliminary report; a full independent review report will follow if required. Completeness is secondary to speed at this stage: get the known facts to the board quickly.

Report Date	
Incident Date / Time	
Reported by	
Board/Committee Notified	

1. Known facts at time of notification

Item	Details known at this time
AI tool involved (name, vendor, tier classification)	
Nature of the incident	
Patient harm or near-miss? (if applicable)	
Number of patients potentially affected	
PHIPA breach threshold triggered?	
Regulatory notification required?	
Financial exposure (estimated)	
Media attention likely?	

2. Immediate actions taken

Describe what management has done since becoming aware of the incident: containment actions, notifications, suspension of AI tool, communications to affected parties.

3. Board direction requested

State clearly what decisions the board is being asked to make at this time, if applicable, including: (a) suspend/continue operation of the AI tool (b) commission independent review (c) authorize regulatory notification (d) approve external communications.

4. Next update

Date and nature of next management update to the board:	

Appendix 5 to sample board AI governance policy:

Management shadow AI detection and response framework

This is a sample response framework for illustrative purposes. This schedule is a sample management-level operational reference for the CIO/CTO, privacy officer, and AMC. It operationalizes the prohibition in section 6 of the Sample Board AI Governance Policy and the IRM sub-categories. Hospitals should seek legal advice prior to adopting this response framework.

1. Definition and scope

Shadow AI (Shadow AI) means any AI system used for hospital work not listed on the hospital's approved AI tools list. This includes personal consumer AI accounts (ChatGPT, Claude, Gemini, Microsoft Copilot, Perplexity, and equivalents); AI embedded in personal productivity applications; departmentally-procured AI tools not reviewed through hospital governance processes; and AI accessed via personal devices. The root cause is structural and includes insufficient approved tooling, AI literacy gaps, and absence of a communicated approved tools list. The governance response must address all three gaps; not merely prohibit the behaviour.

2. Risk categories by role

Highest risk: clinical staff

Physicians, nurses, and allied health professionals using personal LLMs to draft clinical notes, summarize patient histories, or generate care plans are creating risk. Each instance is simultaneously a PHIPA breach and potentially an EDSTA accountability failure. To mitigate this risk, hospitals can deploy an institutionally-approved AI scribe and other AI-enabled tools to assist with these tasks.

High volume: administrative staff

HR staff processing compensation data, finance staff summarizing budget documents and procurement staff drafting vendor communications pose a risk that is lower in individual severity than clinical shadow AI use. However, shadow AI use in administrative roles may occur in much higher volume. To mitigate this risk, hospitals can deploy approved AI productivity tools and clear policy communication.

Moderate, unique exposure: diagnostic and technical staff

Risks include, for example, imaging technologists or pathologists supplementing approved workflows with consumer AI and IT staff using unapproved AI for code generation or system configuration. Mixing approved and unapproved AI in diagnostic workflows creates accountability gaps difficult to reconstruct. To mitigate this risk, the hospital can create an explicit policy that approved workflow tools may not be supplemented with unapproved tools in clinical contexts.

3. The five prevention controls

Control 1: approved tooling

The most effective prevention for clinical shadow AI is an AI scribe that has been approved by the hospital to ensure that PHIPA obligations are complied with. Staff with a usable approved alternative would not need to resort to personal accounts. The CIO/CTO's procurement roadmap must prioritize approved tooling for the highest-volume shadow AI use cases first.

Control 2: published approved AI tools list

The CIO/CTO shall maintain and publish a current, accessible list of all AI tools approved for hospital work, with approved use cases and conditions. Updated within 30 days of any change. Accessible on the hospital intranet.

Control 3: AI literacy training

All-staff training must address shadow AI explicitly: what it is, why it creates legal and patient safety risk, what approved alternatives exist, and how to report an instance. Role-specific modules for clinical, administrative, and IT staff are more effective than generic all-staff modules.

Control 4: network and access controls

The CIO/CTO should assess whether network-level monitoring of consumer AI platform traffic on hospital networks, web content filtering, and endpoint monitoring are appropriate. Technical controls are secondary to approved tooling: staff without approved alternatives will find workarounds. Deploy controls in conjunction with tooling, not as a substitute.

Control 4 should include data-loss-prevention monitoring of content uploaded by staff to consumer AI sites, not only monitoring of which sites are accessed.

Control 5: network-level blocking of high-risk AI sSites

The CIO/CTO shall maintain a list of consumer AI sites blocked at the hospital's network perimeter, reviewed quarterly. Blocking decisions must balance the operational reality that staff will find workarounds where approved tooling is absent (see Control 1), but unsanctioned consumer AI sites with no enterprise data protection feature, no Canadian data residency commitment, or no published security posture should be blocked by default. Detection of unsafe sites by the hospital without blocking is not a sufficient control.

4. Detection and response protocol

The following chart sets out the protocol for detecting and responding to the use of shadow AI. Responses escalate across four levels based on whether patient data was involved and whether the conduct was a first instance, repeated, or deliberate. The framework prioritizes voluntary self-reporting and education over punitive measures, while ensuring that incidents involving patient data trigger the organization's obligations under the PHIPA.

Stage	Response
Detection	Voluntary self-reporting (non-punitive pathway to privacy officer or IT contact); IT help desk incident tracking; quarterly network log review for consumer AI platform traffic; and annual procurement review for departmental AI subscriptions outside the governance process.
Level 1: No patient data, first instance	Education and policy reminder; manager notification; no disciplinary record; and provide information on approved alternatives. Note: This is the default response for most shadow AI incidents, absent aggravating factors.
Level 2: Patient data involved, first instance	Immediate PHIPA breach assessment by privacy officer; breach notification assessment under PHIPA; manager notification; education and remediation; and documentation in breach log. Note: Hospital privacy policy to govern investigation and outcomes.
Level 3: Repeat or deliberate	HR process engagement; formal documentation; escalation to department head and privacy officer; PHIPA breach assessment if patient data involved; and hospital privacy policy to govern investigation and outcomes.
Level 4: Large-scale or systemic	Shadow AI use involving a large number of staff, significant patient data volume, or material regulatory breach requires: • Board notification within 48 hours in accordance with Sample Board AI Governance Policy • Hospital privacy policy to govern investigation and outcomes.

The single most powerful board question on shadow AI:
"Is there a hospital approved, ambient AI documentation tool available to our physicians and nurses today? If not, what is the procurement timeline?"

Appendix 6 to sample board AI governance policy:

AI management committee: terms of reference

The AMC is the institution's central management-level AI decision body, executing the intake, triage, assessment, and approval functions in the stage-gate lifecycle. Every Ontario hospital should establish this committee or designate an existing body with an equivalent mandate. We have set out below an example of the minimum required terms of reference for an AMC.

(a) Responsibilities of the AMC

The AMC reports to the CEO. It supports both stream committees through the CEO, but does not displace them. The AMC is responsible for:

- maintaining the current AI inventory
- operating the stage-gate intake, triage, assessment, and approval process
- assessing all proposed stream 1 Tier 2 AI risk and all stream 2 deployments against the evaluation rubric
- recommending approvals, conditions, and retirements through the CEO to the Audit/Finance Committee (stream 1) or Quality Committee (stream 2) as applicable
- monitoring deployed AI tools against approved performance thresholds, including the four post-go-live failure modes (vendor model changes, workflow drift, weak monitoring ownership, use case scope creep)
- overseeing the Cyber Security Regulation obligations (primary contact designation, CMA, and 72-hour incident reporting)
- conducting assessments of potential unintended consequences of each deployment before approval, including how human-AI interaction patterns may change workflows, role responsibilities, and the amount of time clinicians spend interfacing with technology rather than patients, and
- reporting to the CEO on a schedule the CEO sets, with onward flow to each stream committee.

(b) Composition of AMC

The AMC should include, at minimum, the following expertise and roles:

- Technology (CIO/CTO: chair, or co-chair with senior medical officer for stream 2)
- Clinical (the senior medical officer (VPM or CMIO), acting as co-chair for stream 2 matters)
- Privacy (chief privacy officer)
- Legal (legal counsel)
- Ethics (bioethicist, clinical ethicist, or designated ethics advisor)
- Risk Management (risk management designate)
- Finance (CFO)

- Human resources (VP human resources or delegate)
- Frontline clinical (MAC chair/a designate from the MAC nominated by the chief of staff, plus one frontline nurse or allied health professional), and
- Technology/data (AI or informatics specialist).

Privacy, ethics, and risk management representation are non-negotiable regardless of hospital size. For smaller hospitals, the non-negotiable core is: the senior medical officer, a MAC designate, the chief privacy officer, the CIO/CTO, an ethics advisor, and a risk management designate. External advisors could supplement the core at any hospital size.

The Cyber Security Regulation designated primary contact must be an AMC member. The CIO/CTO is the recommended designate. When the EDSTA-enabling regulations for designated organizations are issued, additional designations or AMC representation may be required and the hospital must update the AMC composition accordingly. The three-line governance model applies: AMC as second-line challenge function; use-case owners in the first line; internal audit or an external advisor in the third line.

(c) Recommendations of the AMC

The AMC makes recommendations to the board and approving authorities, where applicable. It does not itself approve AI deployments. All recommendations are made through the CEO to the relevant approving authority, as follows:

- **Tier 1 AI risk (stream 1)** AMC recommends to CIO/CTO for approval in their executive capacity; Tier 2 AI risk (stream 1): AMC recommends to CIO/CTO with CFO concurrence
- **Tier 2 AI risk (stream 2)** AMC recommends to the senior medical officer for endorsement and to the CEO for approval
- **Tier 3 AI risk** AMC recommends to the CEO for approval with mandatory senior medical officer endorsement and privacy officer and legal counsel review, and
- **Tier 4 AI risk** AMC recommends to CEO, who brings forward to the relevant stream board committee for board resolution.

All AMC recommendations must be accompanied by a documented assessment record. No AI deployment should proceed without approval from the authority designated for its tier. The AMC must meet at minimum quarterly. Called meetings are available when a time-sensitive Tier 2 AI risk or above deployment requires assessment. Quorum rules should be established, which could require the CIO/CTO or designate, the privacy officer or designate and, at minimum, one clinical member.

(d) Conflict of interest

Members must disclose any financial interest in, or advisory relationship with, any AI vendor whose product is under review. Disclosed conflicts trigger recusal from the relevant recommendation. Disclosures and recusals must be documented in AMC minutes.

Schedule E:

Board governance checklists

This is a sample board governance checklist for illustrative purposes. Hospitals should seek legal advice prior to adopting this checklist.

Personal director accountability self-assessment

The following questions are intended to assist board directors in assessing whether they have the requisite insights to discharge their fiduciary duties in respect to AI governance. The questions are adapted from U.S. Senator Ron Wyden's proposed corporate accountability framework,[119] which asks whether each director can personally attest that they exercised reasonable oversight.

Note: If a director cannot affirmatively check each item below, the Board has a gap in its AI governance.

- ☐ I understand my statutory obligations as a board director as they relate to AI governance, including the fiduciary duty, duty of care, duty to monitor, and the proportionality principle governing implementation depth.
- ☐ I understand the board's governance architecture, systems and controls in place for AI deployment, monitoring and incident response.
- ☐ I can identify which board members have AI or cybersecurity expertise or confirm that a formal direct reporting relationship between the CIO or CTO and the relevant board committee is in place, operating, and mandated to surface problems, not just report progress.
- ☐ I can generally describe the structured AI governance report the board received in the last six months and what the minutes recorded.

Board AI governance checklist

The following checklist is intended to assist boards in assessing whether their organization has established the foundational elements of responsible artificial intelligence (AI) governance. It identifies the 20 key questions a board should be able to answer as AI adoption accelerates across clinical, operational, and administrative functions. The questions are ordered from the highest level of governance — the board's own instruments, records and composition — through to regulatory compliance, clinical safety, and operational controls.

119 Hacking America's Health Care: Assessing the Change Healthcare Cyber Attack and What's Next, Senator Ron Wyden, Opening Statement, United States Senate Finance Committee, (May 1, 2024).

Note: A "No" or "Unknown" answer identifies a governance gap requiring immediate attention. Complete annually. Document completion of this checklist in board records.

Yes	No	Unknown	
☐	☐	☐	1. Board has adopted a formal, board-approved AI governance policy with defined materiality thresholds.
☐	☐	☐	2. The governance framework is proportionate to the institution's AI deployment profile.
☐	☐	☐	3. Board minutes document AI governance discussions consistent with business judgment rule process expectations.
☐	☐	☐	4. Board has received the annual AI performance and compliance report, and recorded such receipt in the minutes.
☐	☐	☐	5. Board has reviewed the current AI inventory in the past 12 months.
☐	☐	☐	6. Board has verified that every AI tool in the inventory is classified by AI risk tier (Tier 1-4) and approved at the level the policy requires.
☐	☐	☐	7. Board has verified that AI appears as named risk sub-categories within existing domains of the integrated risk management (IRM) framework.
☐	☐	☐	8. Board skills matrix identifies AI literacy and cybersecurity governance as sought competencies.
☐	☐	☐	9. Either the board has a member with AI or cybersecurity expertise, or a formal direct reporting relationship has been established from the CIO or CTO to the committee with AI governance oversight responsibility.
☐	☐	☐	10. Board AI education has been provided in the past 12 months, documented in minutes.
☐	☐	☐	11. C-suite AI fluency is documented and proportionate to each executive's governance responsibilities.
☐	☐	☐	12. Formal institutional protocols exist for the appropriate selection, implementation, training, maintenance and inspection of AI products.

Yes	No	Unknown	
☐	☐	☐	13. Management has systems in place to verify that AI tools qualifying as software as a medical device (SaMD) that require a medical device licence hold a valid Health Canada licence; not relying solely on vendor representations.
☐	☐	☐	14. All AI systems processing patient data have been assessed for compliance with the *Personal Health Information Protection Act, 2004* (PHIPA), including privacy impact assessments.
☐	☐	☐	15. Ambient AI scribe deployments satisfy the Information and Privacy Commissioner of Ontario's AI scribe guidance and CPSO requirements.
☐	☐	☐	16. Hospital has an *Enhancing Digital Security and Trust Act, 2024* compliance readiness plan; pursuant to the Cyber Security Regulation, a cyber security maturity assessment is scheduled for completion by July 1, 2027 and biennially thereafter, and a designated cybersecurity primary contact and alternate have been identified and provided to the Ministry CISO. The AMC has reported on cybersecurity alignment in at least one of the last two semi-annual reports.
☐	☐	☐	17. Mandatory human verification protocols exist for all AI clinical outputs, and clinician training on automation bias is documented.
☐	☐	☐	18. Continuous post-deployment performance monitoring is in place for all clinical AI tools, with defined performance thresholds and a documented retraining or retirement protocol, and an AI-specific incident reporting protocol is operational and has been tested.
☐	☐	☐	19. All material AI vendor contracts have been reviewed against the institution's key contractual considerations for AI deployments, with depth calibrated to the risk profile of each deployment.
☐	☐	☐	20. The hospital maintains a published approved AI tools list reviewed every six months, an employee AI use policy prohibiting non-approved tools for hospital work, and a current policy governing personal-device access to the EMR and hospital meeting platforms, aligned with IPC guidance and Ontario Health Cyber Security Centre advisories.

About Osler, Hoskin & Harcourt LLP

Osler is a leading law firm with a singular focus — your business. From Toronto, Montréal, Calgary, Vancouver, Ottawa and New York, we advise our Canadian, U.S. and international clients on an array of domestic and cross-border legal issues. Our collaborative "one firm" approach draws on the expertise of over 600 lawyers to provide responsive, proactive and practical legal solutions driven by your business needs. For more than 160 years, we've built a reputation for solving problems, removing obstacles, and providing the answers you need, when you need them.

Osler, Hoskin & Harcourt LLP
Toronto Montréal Calgary Vancouver Ottawa New York | **osler.com**

OSLER