

## **MEMORANDUM**

**TO:** British Columbia Association of School Business Officials

**FROM:** Koffman Kalef LLP

**RE:** Governing Use of Artificial Intelligence

**DATE:** May 21, 2026

---

### **1. Introduction**

From payroll automation and procurement analytics to human resource screening tools and facilities management software, Artificial Intelligence (“AI”) is being embedded in the administrative infrastructure of workplace across British Columbia. The question facing school business officials is no longer whether to use AI, but how to govern its use responsibly, lawfully, and in a manner that upholds the public trust that boards of education are obligated to maintain.

This paper is organized around four themes: (i) the British Columbia legal landscape governing internal AI use; (ii) the specific legal risks and implications for boards of education operations; (iii) the most common governance pitfalls; and (iv) practical best practices for building an AI governance framework.

### **2. Potential for AI Use in Board of Education’s Operations**

#### **2.1 *What is Internal AI Use?***

When we speak of AI governance for internal use, we are not primarily concerned with what students do in classrooms. We are concerned with what boards of education do as employers, procurers, financial managers, and data custodians. Internal AI use includes any system, tool, or software that relies on machine learning, predictive analytics, large language models, or automated decision-making to support administrative or operational functions.

Some examples where board of education’s may encounter AI today include: AI-assisted procurement platforms that score vendor bids or flag anomalies in contract compliance; human resource tools that screen teacher and support staff applications; payroll systems that use predictive modelling for workforce planning or overtime forecasting; facilities management software that uses sensor data and AI to optimize energy use or predict maintenance needs; financial analytics tools that identify budget variances or project grant utilization; and general-purpose generative AI platforms (such as Microsoft Copilot, integrated into the Microsoft 365 environment) used by administrative staff for drafting, summarizing, and research.

The breadth of this list illustrates why a narrow, classroom-focused AI policy is insufficient. The governance gap for administrative AI use is where most legal and operational risk currently lives.

## **2.2 Pace of Adoption vs. Pace of Governance**

Across British Columbia, AI adoption in business offices is accelerating faster than policy frameworks can respond. Vendors are marketing AI-enhanced features as standard upgrades to existing software contracts.

This gap between adoption and governance is the central challenge. It is not unique to British Columbia, it is a national and international phenomenon. But British Columbia's legal environment, anchored in the *Freedom of Information and Protection of Privacy Act* (British Columbia) creates specific obligations that make the governance lag particularly consequential for public bodies.

## **3. Legal Landscape**

### **3.1 Freedom of Information and Protection of Privacy Act (British Columbia)**

The *Freedom of Information and Protection of Privacy Act* (British Columbia) ("**FOIPPA**") is the primary legal framework governing how boards of education, as public bodies, collect, use, disclose, and retain personal information. *FOIPPA* requires school boards to collect personal information only for purposes directly related to their programs or activities, to use it only for the purpose it was collected or for a consistent purpose, and to protect it with appropriate safeguards. Critically, *FOIPPA* also requires that personal information typically be stored and accessed only within Canada, unless a risk assessment is completed to ensure that data is adequately secured.

*FOIPPA* was not written with artificial intelligence in mind. The legislation does not define AI, does not regulate algorithmic decision-making, and does not create explicit obligations around the use of automated systems to process personal data. This is both the central legal challenge and an important governance opportunity. In the absence of prescriptive AI-specific rules in British Columbia legislation, it falls to school boards to interpret and apply *FOIPPA*'s general principles in the AI context.

Significant amendments to *FOIPPA* came into force beginning February 1, 2023, under Bill 22. These included mandatory Privacy Management Programs, which required each public body to designate a privacy officer, maintain documented privacy policies, conduct Privacy Impact Assessments (PIAs) and maintain breach-notification procedures.

### **3.2 Privacy Impact Assessments and AI**

A Privacy Impact Assessment is a structured process for identifying the privacy risks of a program or system and documenting how those risks will be mitigated. Amendments to *FOIPPA* made PIAs mandatory for new programs involving personal information. AI systems almost invariably involve personal information such as employee records, student data, contractor information, or financial records linked to individuals.

Conducting a PIA for an AI system is more complex than for a conventional software tool, for several reasons. First, AI systems often cannot fully explain how they reach conclusions. This makes it difficult to document the decision making logic, as required in a thorough PIA. Second, AI systems trained on

historical data may encode past patterns of discrimination, creating privacy and human rights risks that are not immediately visible. Third, AI systems often involve data flows to third party vendors, including cloud providers whose servers may be located outside Canada, potentially violating *FOIPPA*'s storage requirements.

### **3.3 Clearview Decision – Precedent for AI Accountability**

In December 2024, the British Columbia Supreme Court issued its decision in *Clearview AI Inc. v. Information and Privacy Commissioner for British Columbia*, 2024 BCSC 2311 (“**Clearview**”). The case arose from Clearview AI's practice of scraping billions of publicly accessible photographs from the internet to build a facial recognition database without obtaining consent from the individuals depicted. BC's Office of the Information and Privacy Commissioner (“**OIPC**”) found this practice violated the *Personal Information Protection Act* (British Columbia) (“**PIPA**”) and issued a binding order. *PIPA* is the privacy legislation applicable to private parties in British Columbia unlike *FOIPPA* which is applicable to public bodies.

Clearview challenged the OIPC's jurisdiction and the order itself. The BC Supreme Court upheld the Commissioner's authority and the order. The significance of this decision for school boards is not that they are at risk of running facial recognition databases. Rather, AI systems that process personal information without meaningful consent and without a clear, legitimate purpose will be found to violate BC privacy law. This principle applies as much to an AI-assisted human resource screening tool as it does to a global facial recognition company.

### **3.4 The BC Human Rights Code and Algorithmic Discrimination**

The BC *Human Rights Code* prohibits discrimination in employment on the basis of protected characteristics including race, gender, age, disability, family status, religion, and sexual orientation. These prohibitions apply fully to AI-assisted employment decisions. A board of education that uses an AI tool to screen applications, rank candidates, assess performance, or inform termination decisions is legally responsible for discriminatory outcomes, even if those outcomes were produced by an algorithm rather than a human being, and even if the discrimination was unintentional.

If an AI screening tool systematically disadvantages candidates from protected groups (for example, because the training data reflected historical hiring patterns that underrepresented certain demographics) the board of education may be found to have discriminated, regardless of intent. BC's Human Rights Tribunal has not yet issued a major decision specifically on AI-driven discrimination, but the legal framework is clear and the risk is real.

## **4. Legal Risks by Operational Domain**

### **4.1 Human Resources and Recruitment**

Human resources is one of the highest-risk domain for internal AI use in school board administration. The intersection of employment law, human rights law, privacy law, and labour relations creates a complex web of obligations that AI tools can either help navigate or inadvertently breach.

AI in recruitment, including resume-screening algorithms, automated interview scoring, and predictive hiring tools, carries substantial legal risk. For boards of education, which have active obligations under

employment equity principles and collective agreements, this risk is particularly acute. A single biased algorithm can produce discriminatory outcomes across hundreds of hiring decisions simultaneously.

#### **4.2 Finance and Procurement**

AI in financial management and procurement is generally lower-risk from a human rights perspective but carries significant concerns related to accountability, transparency, and procurement law. The BC Procurement Policy Framework requires that public body procurement processes be fair, open, and transparent. An AI tool that scores vendor bids, predicts contract compliance, or recommends contract awards must itself be transparent and auditable.

If a vendor challenges a procurement decision that was informed by AI, arguing, for example, that the scoring algorithm was opaque, inconsistent, or biased against certain vendor types, the board of education must be able to explain the decision making process. An AI system that cannot be audited or explained creates legal exposure in procurement disputes that would not exist with conventional evaluation matrices.

Financial analytics AI that uses employee or contractor must comply with *FOIPPA*'s use-limitation principle: data collected for one purpose cannot be used for a different, inconsistent purpose without fresh authority. Using payroll data to train or power an AI system designed for financial modelling may require a new *FOIPPA* analysis and potentially a new collection authority.

#### **4.3 Facilities Management and Operational AI**

AI applications in facilities management including smart building systems, predictive maintenance platforms, and energy optimization tools are generally the lowest-risk category from a personal information perspective, provided they operate on anonymized sensor and equipment data rather than data linked to individuals. However, some operational AI tools do process personal information: security camera systems with AI-enhanced monitoring, occupancy sensors that can be linked to individual movements, or access control systems that track individual entry and exit patterns.

Any AI system that processes data capable of identifying an individual including biometric data, location data, or behavioural patterns could fall within *FOIPPA*'s definition of personal information and must be governed accordingly. School boards should conduct a careful classification exercise for each operational AI tool to determine whether personal information is involved, and if so, what legal authority exists for its collection and use.

#### **4.4 Generative AI in Administrative Work**

Perhaps the most pervasive and least governed category of internal AI use is the deployment of general purpose generative AI tools by administrative staff. These are tools like Microsoft Copilot (integrated into Microsoft 365), Google Gemini, or standalone large language model interfaces. These tools are used by staff across finance, HR, communications, and executive functions to draft documents, summarize reports, generate correspondence, and answer complex questions.

School boards should develop their own internal AI usage policies that define what information may and may not be inputted into AI tools, how AI-generated outputs must be reviewed and verified before use, and what records management obligations apply to AI-assisted correspondence.

## **5. Common Governance Pitfalls**

Understanding where AI governance fails in practice is as important as knowing what good governance looks like.

### **5.1 *Shadow AI: Adoption Without Authorization***

Shadow AI refers to the use of AI tools by employees or departments without formal authorization, procurement review, or privacy assessment. A teacher or administrator who creates a free account on an AI platform and uses it to process student or employee data has created a privacy breach exposure for the board of education, even if the employee had entirely benign intentions.

Shadow AI is particularly difficult to manage because many AI tools are embedded in commonly used consumer applications, available at no cost, and marketed as productivity tools rather than data-processing systems. School boards must communicate clearly to all staff that any tool used to process work-related information including personal information, confidential financial data, or legally privileged communications must be reviewed and approved before use, regardless of whether it is free or commercially available.

### **5.2 *Treating Vendor Assurances as Legal Compliance***

When procuring AI enhanced software, school boards often rely on vendor representations about privacy compliance, data security, and *FOIPPA* compatibility. These representations are not legal compliance. School boards must negotiate specific contractual protections with AI vendors, including: confirmation of data storage locations; prohibition on using school board data to train AI models without explicit consent; breach notification timelines consistent with *FOIPPA* obligations; audit rights; and data deletion obligations on contract termination. These terms should be reviewed by legal counsel before the contract is executed.

### **5.3 *The “Human in the Loop” Fallacy***

A common response to concerns about automated AI decision making is to assert that “a human is always in the loop”. This assertion provides less legal protection than many administrators assume. A human who routinely accepts AI recommendations without meaningful independent review is not truly providing oversight. They are essentially rubber stamping an automated decision. Courts and regulators are alert to this distinction.

For human in the loop oversight to be legally meaningful, the human reviewer must: understand how the AI system reached its recommendation; have access to the underlying data and reasoning; have the genuine authority and practical ability to override the AI recommendation; and document the basis for their decision independently of the AI output. Policies that mandate human review of AI recommendations must therefore also mandate meaningful human capability to exercise that review.

### **5.4 *Failure to Manage AI Vendor Lock-In and Data Portability***

School districts that build operational dependencies on AI vendor platforms may find themselves unable to exit those relationships without losing access to years of operational data. If student, employee, or financial data has been processed through a vendor’s AI system, the board of education

must ensure it can retrieve that data in a usable format when the contract ends and that the vendor's obligations to delete residual data are enforceable. Inadequate attention to data portability and end of contract data handling at the procurement stage can result in long-term exposure and financial risk.

## **6. Best Practices: Building an AI Governance Framework**

An effective AI governance framework for boards of education does not need to be elaborate. It needs to be clear, consistently applied, and legally grounded.

### **6.1 Establish an AI Governance Policy**

Every school board should have a AI Governance Policy that defines the board of education's approach to AI. This policy should establish the principles that will guide AI use (including human accountability, transparency, privacy protection, and equity); define what qualifies as an AI system for governance purposes; establish the approval process for adopting new AI tools; assign accountability for AI governance; and require periodic review as the AI landscape evolves.

The policy need not be technically complex. It should be written in language accessible to non-technical staff and should be explicitly linked to the board of education's existing privacy policy.

### **6.2 Implement an AI Risk Classification System**

Not all AI tools carry the same level of risk. A building energy optimization system that processes no personal data is categorically different from an AI system that scores employee performance or screens job applications. School boards should implement a tiered risk classification system that matches the level of review and approval required to the level of risk presented by each AI application.

- Tier 1 (Low Risk): AI tools that process no personal information and have no direct impact on decisions affecting individuals. Examples: predictive maintenance tools using equipment sensor data; AI-assisted document formatting tools.
- Tier 2 (Moderate Risk): AI tools that process personal information but are not used to make or inform significant decisions about individuals. Examples: AI-assisted scheduling tools; generative AI used for administrative drafting with appropriate input controls.
- Tier 3 (High Risk): AI tools that process personal information and are used to inform significant decisions about individuals, including employment, financial, or disciplinary decisions. Examples: AI hiring tools; AI performance management systems; AI budget forecasting tools using personnel data.

### **6.3 Require Privacy Impact Assessments for All Tier 2 and 3 AI Tools**

PIAs should be undertaken for new programs involving personal information. For AI tools, the PIA must specifically address: what personal information is collected and from what sources; the legal authority for collection; how and where data is stored ; how the AI uses the data, including whether it is used for model training; what the potential privacy risks are; what safeguards are in place; how individuals can exercise their rights (including access and correction); and how the board of education will respond to a breach.

The PIA should be completed before the AI tool is deployed and not after. It should be a living document updated whenever the tool or its data practices change materially.

#### **6.4     *Build AI Obligations into Contracts***

Every contract for an AI enhanced tool or service should include specific AI governance provisions. Working with legal counsel, school boards should ensure vendor contracts for AI tools or services cover the following:

- Data location: Explicit commitment that all personal information will be stored and processed in Canada, or identification of specific exceptions with school board consent.
- No model training on school board data: Clear prohibition on using school board personal information to train or improve AI models without explicit written consent from the school board.
- Transparency and explainability: Commitment to provide, on request, a plain-language explanation of how the AI system reaches its outputs or recommendations.
- Audit rights: Right to audit vendor compliance with privacy and AI governance obligations, including right to review bias testing results.
- Breach notification: Obligation to notify the school board within a specified time of any actual or suspected breach involving school board data.
- Data deletion: Obligation to delete all school board data within a specified period after contract termination, with written certification of deletion.
- Liability allocation: Clear statement of each party's liability for claims arising from AI-generated errors, bias, or privacy breaches.

#### **6.5     *Establish a Staff AI Usage Policy***

School boards must communicate clearly to all employees what they may and may not do with AI tools in the course of their work. An employee AI Usage Policy should define approved AI tools and the contexts in which they may be used, prohibit the input of personal information, confidential information, or legally privileged information into non-approved AI tools, require that all AI-generated outputs be reviewed and verified by a human before use or dissemination, and establish records management requirements for AI-assisted correspondence and decision-making.

This policy should be delivered through mandatory training, not merely posted on an intranet. Staff should understand not only what the rules are, but why they exist.

#### **6.6     *Designate an AI Governance Lead***

AI governance cannot be an afterthought or a shared responsibility with no clear owner. Each school board should designate an AI Governance Lead who is responsible for maintaining the AI tool inventory, coordinating PIAs, reviewing vendor contracts, responding to AI-related incidents, and keeping the

board and senior leadership informed of emerging legal developments. This role should have the authority to pause or prohibit AI tool deployments that do not meet the district's governance standards.