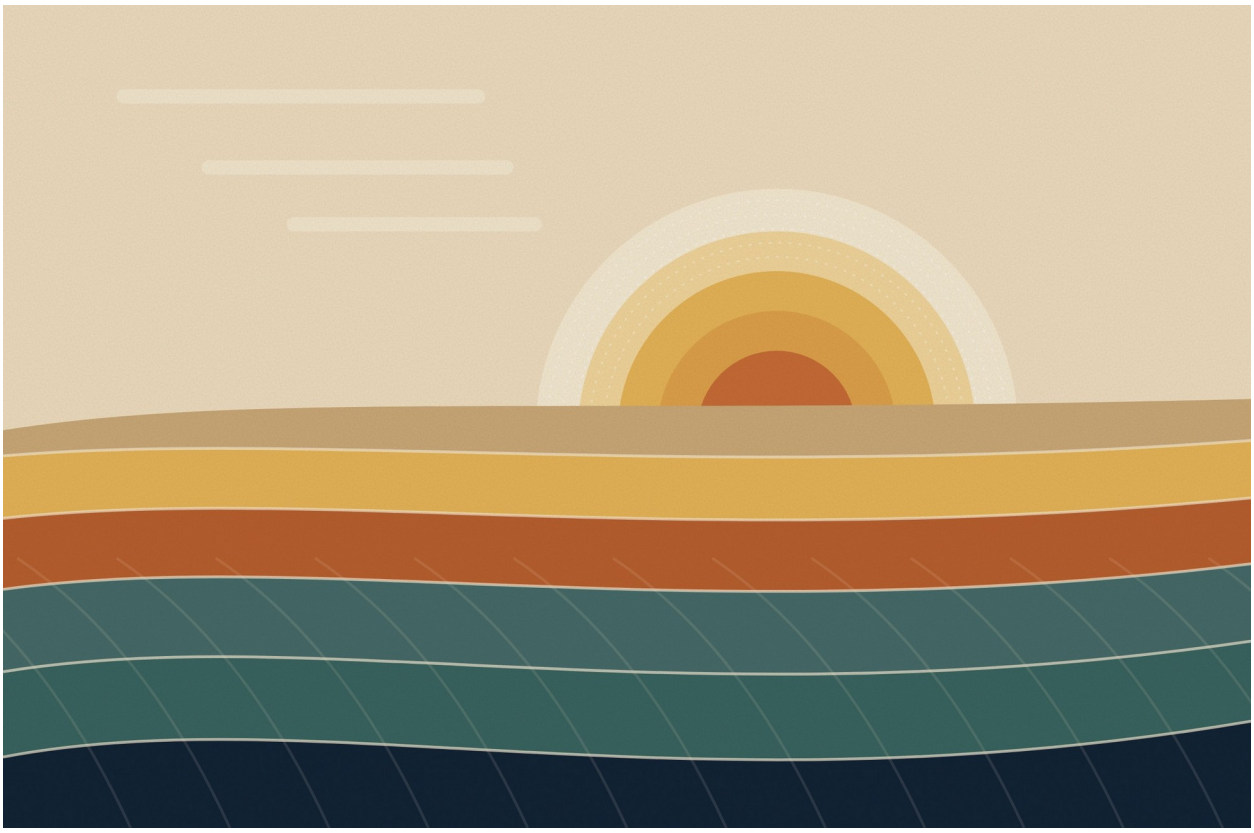


Artificial Intelligence in Education

A Policy Advisory for Saskatchewan Boards of Education

Prepared by Dr. Alec Couros, Faculty of Education, University of Regina
Informed by the SSBA Artificial Intelligence in Education Working Advisory Group
Directed by Resolution AGM 25-04, adopted at the 2025 Annual General Meeting

Version 1.0 · 12 October 2026

SASKATCHEWAN School Boards ASSOCIATION

Foreword

DR. TED AMENDT



The Saskatchewan School Boards Association's Artificial Intelligence in Education Working Advisory Group (SSBA AI WAG) was established following a resolution brought forward by our members at the 2025 Fall General Assembly. The SSBA moved quickly on this important issue identified by our members and brought together sector partners and subject matter experts for a robust process of discussion, consideration, and direction-setting that has resulted in this document and the associated tools that will be valuable to member boards of education.

Acknowledgements are in order, as the SSBA thanks the many members, seventeen in total, who were appointed to the AI WAG.

First and foremost, the SSBA thanks Dr. Alec Couros, who served on the SSBA AI WAG as a member representing the University of Regina and worked diligently as our lead consultant and developer to guide the process, shape this document, and build the fantastic tools associated with it. His knowledge and leadership have been essential to this process and to laying the foundation for future work in the sector.

Next, thank you to the co-chairs of the SSBA AI WAG. Jerome Niezgoda, SSBA Executive Catholic Constituency Representative, and Lisa Grudnizki, SSBA Executive Southern Constituency Representative, have both dived deeply into this subject and helped move this work forward.

Our co-chairs were supported by additional SSBA appointees bringing important voices to bear. Anne-Marie Rollo, from Saskatoon Public Schools, represented the perspective of a board member and an education expert. Simone Gessler, from Greater Saskatoon Catholic Schools, modelled collaboration by helping ensure the WAG's work aligned with operational efforts through the Saskatchewan Artificial Intelligence Leadership Network (AILN), which she co-leads. Simon Bird brought his perspective as an educational leader and participant in the SSBA's Indigenous Educators Summit on Artificial Intelligence, held as part of this process and bringing together voices from across the province.

We were fortunate to have significant partnership representation from the Government of Saskatchewan, including Legislative Secretary for Education Barret Kropf; Ministry of Education representatives Jeff Bassendowski and Megan Cottenie; and, from the Saskatchewan Distance Learning Corporation, Kelli Boklaschuk, who also brought perspective as the other co-lead of the AILN.

From our sector partners, participation in the WAG was significant, and we thank the following individuals for their efforts. Derek Barss and Ted Zurakowski represented the Saskatchewan Teachers' Federation. Nigel McCarthy (Conseil des écoles fransaskoises) and Stacy Lair (North East School Division) represented the League of Educational Administrators, Directors and

Superintendents of Saskatchewan. The Saskatchewan Association of School Business Officials was represented by Jesse Green and Justin Arendt (both of Horizon School Division). The University of Saskatchewan was represented by Dr. Paula MacDowell.

The work of the AI WAG was also supported by numerous SSBA staff members and service areas.

This document reaches boards in a timely way as artificial intelligence continues to evolve and remains an important issue with implications for policy and operations. The guidance provided here isn't intended to be prescriptive, but to empower boards with information and education for the considerations they might face in the future.

With longer-term coordination in the sector anticipated, including through the AILN at the operational level, our member boards are now well equipped through the results of this work. We recognize the importance of continued collaborative leadership in this area to ensure our students, staff, families, and communities are prepared to navigate the future with clarity, consistency, and protection.

Statement on the Use of Artificial Intelligence

In keeping with principles of transparency, the authors and the SSBA disclose the role of generative AI in the development of this work. Generative AI tools supported aspects of the research, brainstorming, organization, and writing refinement, as well as a portion of the technical development and organization of the companion website. Drawing on consultation with the SSBA Artificial Intelligence in Education Working Advisory Group, a survey of its members and partners, and perspectives shared by participants at the Indigenous Educators Summit described in Section 1, the authors drafted the document's substantive content, positions, and recommendations. Information generated or identified with AI assistance was verified against original sources, and the document was reviewed by both the Working Advisory Group and the SSBA.

Licensing

© 2026 Saskatchewan School Boards Association. Except where otherwise noted, this work is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International Licence (CC BY-NC-SA 4.0). You may share and adapt the material for non-commercial purposes provided that you give appropriate credit, provide a link to the licence, indicate whether changes were made, and distribute adaptations under the same or a compatible licence. This licence does not apply to third-party material, quotations, logos, or trademarks. Use or adaptation of the material does not imply endorsement by the Saskatchewan School Boards Association.

Additional permission for sample policy statements. Boards of education and other public or non-profit educational authorities may reproduce and adapt the sample policy statements for inclusion in their own policies and administrative procedures without applying the ShareAlike condition to the resulting policy or procedure as a whole. Attribution to this Advisory remains required. Full licensing information: https://saskschoolboards.ca/ai/licensing.html

OCAP® is a registered trademark of the First Nations Information Governance Centre (FNIGC). For more information, see https://fnigc.ca/ocap-training/.

Legal information. This Advisory and its companion resources provide general information to support policy development. They do not constitute legal advice from the SSBA or its legal counsel and are not a substitute for advice about a board's particular circumstances. Boards should obtain legal advice when interpreting legal requirements or developing and applying local policies. The guidance does not replace obligations under applicable law.

Contents

Foreword 2
Statement on the Use of Artificial Intelligence 4
Licensing 4
Preface 7
Executive Summary 8
What the Advisory provides 8
How the Advisory should be used 10
1. The SSBA and the Origin of This Policy Advisory 11
1.1 The Saskatchewan School Boards Association 11
1.2 Resolution AGM 25-04 11
1.3 The AI Working Advisory Group 11
1.4 From draft to adoption 12
2. Scope of This Policy Advisory 13
2.1 A governance document 13
2.2 Spheres of responsibility: govern, influence, monitor 13
3. The Context for This Policy Advisory 15
3.1 From generative AI to embedded infrastructure 15
3.2 Understanding the broader context 15
3.3 The Saskatchewan context 16
4. Guiding Principles 18
4.1 Human-centred education 18
4.2 Student well-being and children's rights 18
4.3 Privacy, security, and responsible data use 19
4.4 Equity, inclusion, and access 19
4.5 Ethical and responsible use 19
4.6 Teacher professional agency 20
4.7 Indigenous authority and self-determination 21
4.8 Transparency and accountability 21
4.9 Adaptability and continuous learning 22
4.10 How the guiding principles function throughout this document 22
5. Indigenous Authority, Knowledge, and AI Governance 24
5.1 Direction on Indigenous language, ceremony, and sacred knowledge 24
5.2 Accuracy does not establish authority 25
5.3 Community knowledge and data in division custody 26
5.4 Decisions affecting students and staff 26
5.5 Engagement and continuing dialogue 27
6. Policy Domain Guidance 28
6.1 AI literacy, professional learning, and well-being 29
6.2 Academic integrity and assessment 32
6.3 Privacy, security, and data governance 35
6.4 AI tools, procurement, and intellectual property 38
6.5 Prohibited and restricted uses 41
6.6 Equity, accountability, and community engagement 47
6.7 Applying the domains: division administration and business functions 50
7. Learning from Other Jurisdictions 56
Patterns across jurisdictions 56

8. Keeping This Policy Advisory Current 57
References 58
Appendix A: Board Discussion Questions and Sample Policy Statements 62
A.1 AI literacy, professional learning, and well-being 63
A.2 Academic integrity and assessment 65
A.3 Privacy, security, and data governance 68
A.4 AI tools, procurement, and intellectual property 71
A.5 Prohibited and restricted uses 74
A.6 Equity, accountability, and community engagement 78
A.7 Applying the domains: division administration and business functions 81
Appendix B: Board AI Governance Self-Check 86
Using the self-check 87
Part One: Required safeguards 88
Part Two: Governance maturity 90
Recording and interpreting the results 101
Domain response profile 102
Board discussion and action planning 103
Appendix C: Understanding AI: Plain-Language Concepts 104

Preface

Artificial intelligence (AI) has become part of everyday life in Saskatchewan schools. AI is already embedded in the software students use to write, the platforms staff rely on, and the systems school divisions use to administer programs, communicate, and report. At the 2025 Annual General Meeting, member boards responded to this reality by directing the Association, through Resolution AGM 25-04, to develop a comprehensive policy advisory document on artificial intelligence in education. This document is the result of that direction.

This Advisory is intended to support boards of education in developing their own AI-related policies. It provides nine guiding principles; accessible background information, including plain-language AI concepts; guidance on Indigenous authority, knowledge, and AI governance; policy guidance for six key domains; a bank of sample policy statements (Appendix A); and an optional Board AI Governance Self-Check with priority action planning (Appendix B). The companion website provides the AI Literacy and Use Continuum, an illustrative, grade-banded governance reference, and detailed annotated examples from other jurisdictions. Nothing in this Advisory is mandatory; boards may adopt, adapt, or choose not to use any part of this guidance. Additionally, nothing in this document displaces the authority of a board of education or substitutes for the professional judgement of educators within their respective roles.

This Advisory is structured around three deliberate boundaries. First, it is not intended as an implementation manual and does not prescribe operational procedures. The work of vetting tools, delivering professional learning, and supporting classrooms remains with each school division and its administration. Second, it does not set out curriculum or direct classroom instruction. Decisions about what students learn are made through provincial processes, and references to curriculum and learning in this document are included only to support informed governance. Third, it recognizes that boards cannot control the broader forces beyond their authority, including the direction of the AI industry and the technological, social, educational, and economic changes associated with AI. It does, however, provide the relevant background information that boards require to understand these developments, consider their implications for students and school communities, and respond through responsible governance. Where coordinated action may be of benefit, the document identifies selected opportunities for sector collaboration and advocacy through the Association.

Finally, the entire document rests on one conviction: education is a human undertaking, and AI belongs in schools only when it supports student learning, strengthens professional judgement, and sustains public trust. In Saskatchewan, this human-centred commitment carries an ethical and relational responsibility to respect Indigenous self-determination and recognize the authority of First Nations and Métis communities over how their languages, ceremonies, sacred knowledge, and data are represented, used, or reproduced through AI systems, as discussed in Section 5.

AI is evolving more quickly than traditional policy cycles, so the Advisory is grounded in durable principles rather than in the details of particular tools. The SSBA will update the Advisory as needed to reflect changes in technology, law, evidence, and sector experience. A companion website has been developed to house the material that requires more frequent updating: it holds the current edition of the Advisory, templates and practical supports, and current information, resources, and learning opportunities as they emerge. Section 2.2 describes the site in full.

Executive Summary

AI is already embedded in many instructional and administrative systems used by school divisions, and its capabilities, terms, and risks continue to change. A simple decision to permit or prohibit AI is therefore insufficient. Boards need durable principles, clear boundaries for higher-risk uses, and administrative processes for approving and reviewing tools, proposed uses, and higher-impact technology-enabled practices, together with expectations for privacy, professional learning, oversight, and continuing review.

Recognizing this need, Saskatchewan's boards of education directed the development of this Advisory. At the Saskatchewan School Boards Association's 2025 Annual General Meeting, member boards adopted Resolution AGM 25-04, calling for a comprehensive policy advisory document on artificial intelligence in education. The resolution called for guidance on ethical and effective AI use; attention to equity, access, and student data privacy; a framework that boards could adapt to their local contexts; and a review process to keep the guidance current as technologies evolve (SSBA, 2025a).

To fulfil the direction, the Association established the Artificial Intelligence in Education Working Advisory Group, surveyed members and partners, and convened an Indigenous Educators Summit in June 2026. The Advisory translates this work into guiding principles, policy domains, and practical supports for Saskatchewan boards.

What the Advisory provides

Nine guiding principles: Drawing on these consultations and a review of Canadian and international policy, guidance, and research, Section 4 sets out nine guiding principles for local policy development. They are grounded in the conviction that education is a human undertaking and that AI belongs in schools only when it supports student learning, strengthens professional judgement, and sustains public trust. The principles apply across the six policy domains described in Section 6 and are intended to guide decisions about specific tools and practices as technologies continue to change:

1. human-centred education;
2. student well-being and children's rights;
3. privacy, security, and responsible data use;
4. equity, inclusion, and access;
5. ethical and responsible use;
6. teacher professional agency;
7. Indigenous authority and self-determination;
8. transparency and accountability; and
9. adaptability and continuous learning.

Together, these principles support boards in considering not only what AI systems can do, but also how their use affects students, educators, relationships, rights, equity, privacy, Indigenous authority, and public accountability. Protecting students' critical thinking, well-being, and human relationships from over-reliance on AI is a central concern throughout the Advisory. Section 5 expands upon the

principle of Indigenous authority in relation to knowledge, language, and data, and its guidance applies across the policy domains.

Six policy domains: The six policy domains identify the areas in which boards can apply the guiding principles through local policy. For each domain, the Advisory provides background information, questions for board consideration, and sample policy statements that may be selected, combined, and adapted to local circumstances. The six domains are:

1. AI literacy, professional learning, and well-being;
2. academic integrity and assessment;
3. privacy, security, and data governance;
4. AI tools, procurement, and intellectual property;
5. prohibited and restricted uses; and
6. equity, accountability, and community engagement.

Teaching and learning is not presented as a separate domain because it is addressed throughout the Advisory, particularly through AI literacy and well-being in Section 6.1, assessment in Section 6.2, and the AI Literacy and Use Continuum on the companion website. The six domains form the working core of the Advisory. Section 6.7 shows how guidance from these domains applies to division administration and business functions, including human resources, finance, facilities, transportation, records, meetings, and communications.

Clear positions on select key issues: On most matters, the Advisory offers options. On a limited number, it takes a firm position based on legal obligations, rights, safety, and available evidence. For instance, it recommends against procuring AI-text detection tools to determine academic misconduct. AI-detection scores should not be treated as proof of misconduct or as the sole basis for an allegation or decision. It directs attention instead to assessment design, clear task-level expectations, honest disclosure of AI use, and evidence of students' learning processes. Assessment and evaluation of student learning remain the professional responsibility of teachers, and no fully automated system should make a final decision that significantly affects a student, staff member, or other person.

The Advisory also affirms that school divisions' obligations as local authorities under The Local Authority Freedom of Information and Protection of Privacy Act (LA FOIP) apply fully to AI systems. Personal information must not be entered into an AI tool that the division has not assessed and approved for that purpose. Before an external AI service handles personal information, the division should assess the proposed use and ensure that the provider is governed by written agreements and safeguards that meet LA FOIP requirements.

The Advisory also calls for clear local prohibitions on school-connected AI-generated or AI-altered intimate or sexualized images, audio, or video; impersonation intended or reasonably likely to deceive, harass, exploit, defraud, or cause harm; fully automated consequential decisions; and covert or continuous AI-enabled surveillance beyond what law and division policy expressly authorize. Student-facing conversational AI, analytics concerning identifiable students, and other higher-risk uses should proceed only for an approved purpose and with defined safeguards and meaningful human oversight.

Finally, consistent with direction received through the Indigenous Educators Summit, the Advisory states that a division should not use AI to teach an Indigenous language or to generate, reproduce, or prepare material involving ceremony, ceremonial protocol, or sacred knowledge unless the relevant First Nation, Métis government or community, or First Nations education authority has directed or authorized the use. It also calls for respect for Indigenous self-determination and data sovereignty and for ongoing engagement through appropriate relationships and protocols.

Practical tools and companion resources: The Advisory includes jargon-free explanations of the AI concepts it uses (Appendix C); the AI Literacy and Use Continuum, an illustrative Kindergarten to Grade 12 governance reference available through the companion website; a summary of patterns from Canadian and international jurisdictions, with detailed annotated examples to be maintained on the companion website; an optional Board AI Governance Self-Check (Appendix B); and sample policy statements organized by domain and theme (Appendix A).

The companion website (Section 2.2) houses the Advisory and its supporting resources.

Saskatchewan-specific and appropriately bounded guidance: The Advisory is designed for Saskatchewan's varied urban, rural, northern, First Nations, and Métis contexts, including low-bandwidth and non-one-to-one environments. It is built on the assumption that most classroom AI use will occur through teacher- or school-provisioned devices, and it reflects current provincial policy governing cellphone use during instructional time.

The Advisory distinguishes among what a board can govern, what it can influence, and what it must monitor. Boards can govern local policy, expectations, accountability, and reporting. They can influence wider matters through sector collaboration and advocacy. They must also monitor forces beyond their direct control, including developments in the AI industry, commercial dependency, society-wide adoption, the changing legal environment, and the environmental and resource demands of AI systems.

How the Advisory should be used

Readers pressed for time can use the Executive Summary, Section 6, and Appendices A and B as the working core of the document.

This document operates at the level of board governance, providing principles, policy commitments, boundaries, and oversight expectations. The operational work of assessing tools and proposed uses, developing procedures, delivering professional learning, and supporting classrooms remains the responsibility of division administration. Curriculum remains a provincial responsibility. Section 2 develops this distinction further.

A board may find it constructive to complete the self-check in Appendix B, identify the policy domains most urgent in its local context, and adapt the sample policy statements to its own values, responsibilities, and circumstances. Early priorities may include clear privacy expectations and firm prohibitions, communication with staff and families, engagement with First Nations and Métis partners, and an established process for evaluating and approving AI tools and proposed uses.

The SSBA will update the Advisory as needed to reflect changes in technology, law, research, and sector experience. The companion website will provide updates and resources between editions.

1. The SSBA and the Origin of This Policy Advisory

1.1 The Saskatchewan School Boards Association

The Saskatchewan School Boards Association (SSBA) is a non-profit, non-partisan, voluntary membership organization that provides leadership and support for its members as the unified voice of Saskatchewan's 27 democratically elected public, Catholic, and Francophone boards of education (SSBA, 2026). The Association represents boards of education themselves rather than individual board members, and its membership reflects the rural, urban, and northern diversity of the province. Membership is also open to First Nations education authorities (SSBA, 2026).

The SSBA sets its priorities largely through member resolutions at the Annual General Meeting (SSBA, 2026). Its policy advisories support boards in developing sound local policy while respecting each board's autonomy to decide for its community.

1.2 Resolution AGM 25-04

At the 2025 Annual General Meeting, member boards adopted Resolution AGM 25-04, sponsored by South East Cornerstone School Division No. 209 (SSBA, 2025b). The adopted resolution directed the Association to create a comprehensive "Policy Framework Advisory Document on Artificial Intelligence in Education" (SSBA, 2025a). It required that the document include, but not be limited to:

- guiding policy considerations for ethical and effective use of AI in education;
- policy considerations related to equity, access, and student data privacy;
- a framework for member boards to adapt and apply AI policy within their local contexts; and
- a review process so the document remains current as technologies evolve.

1.3 The AI Working Advisory Group

To carry out the resolution, the Association established the Artificial Intelligence in Education Working Advisory Group (WAG), bringing together representatives from school boards, sector partners, and post-secondary institutions. From the outset, the WAG committed to developing Saskatchewan-specific guidance that reflects the province's students, educators, communities, geography, and values. Its role was to provide guidance and shared direction rather than prescriptive rules, recognizing school-division autonomy while promoting greater consistency and shared resources.

The WAG began work in January 2026 and met throughout the year. Two consultation activities primarily shaped the Advisory. A spring 2026 survey of WAG members and partners gathered detailed input on the document's tone, scope, guiding principles, risks, opportunities, and the practical tools boards most needed. The findings are documented separately (Couros, 2026).

The SSBA also convened the Indigenous Educators Summit in Saskatoon on June 4, 2026. First Nations and Métis educators, leaders, and a Knowledge Keeper shared perspectives on AI in education, and their guidance informs Section 5 and other parts of the Advisory. The document also

draws on provincial, national, and international guidance, Canadian board and division policies, and current research cited throughout.

1.4 From draft to adoption

An internal review draft of the Advisory was prepared in August 2026. Member and organizational feedback was gathered through September and October 2026. This version will be presented at the SSBA Annual General Meeting in Regina in November 2026.

2. Scope of This Policy Advisory

2.1 A governance document

This Advisory is written for boards of education. It identifies principles to guide AI governance; policy commitments intended to support responsible use and protect students, staff, information, and public trust; questions boards should consider when establishing and reviewing direction; and matters they should monitor as policy is implemented. It does not prescribe operational procedures, set curriculum, or direct classroom pedagogy. Operational procedures remain the responsibility of division administration, curriculum remains a provincial responsibility, and classroom pedagogy remains a matter of professional judgement within educators' respective roles. To support board governance, the Advisory provides sample policy statements that boards may adapt to their local contexts.

Board policy should create room for purposeful exploration of AI while keeping its use under accountable human direction. Coglianese and Crum (2025) describe a "leash" approach to AI risk governance: controls adapt to the use and context, while people remain responsible for identifying risks, monitoring effects, and intervening when needed. Applied to education, this means supporting approved, developmentally appropriate exploration while maintaining legal obligations, privacy protections, Indigenous authority, and the boundaries for prohibited uses.

2.2 Spheres of responsibility: govern, influence, monitor

AI decisions seldom fall neatly within a single area of responsibility. A decision about licensing an AI tool, for example, may involve board policy and risk; administrative procurement, privacy, security, and support; teaching and learning; and Indigenous data governance. To define the Advisory's scope, this section distinguishes among three spheres of board responsibility: what a board can govern, what it can influence, and what it must monitor.

What a board can govern: Local policy, including guiding principles; expectations for permitted, restricted, and prohibited uses; privacy and responsible data use; policy requirements for approving and reviewing tools, proposed uses, and higher-impact technology-enabled practices; professional learning priorities; accountability; and reporting. Division administration translates these commitments into procedures and practice. This sphere is addressed primarily in Sections 4 through 6 and Appendix A.

What a board can influence: Matters a single board cannot decide but can help shape through sector collaboration and advocacy. Examples include equitable technology funding, provincial support for shared tool evaluation, AI literacy in pre-service teacher education, and provincial or federal legislative direction. Section 6.6 identifies circumstances in which interdivisional collaboration, coordinated procurement, shared services, or advocacy may be appropriate, and Appendix A includes related policy options.

What a board must monitor: Forces beyond its direct control, including changes in the AI industry and legal environment, commercial dependencies, society-wide adoption, and the environmental and resource demands of AI systems. Boards cannot control these developments, but sound governance requires understanding how they may affect students, staff, communities,

finances, and operations. Section 3 provides background, while Sections 6.4, 6.6, 6.7, and 8 address their local implications and continuing review.

THE COMPANION WEBSITE

The companion website is the Advisory's online home. It provides the current edition and updates issued between formal reviews. It also provides the AI Literacy and Use Continuum, an illustrative governance reference that helps boards understand developmental considerations from Kindergarten to Grade 12 without prescribing curriculum or classroom practice. The site extends the document through a guided policy generator drawing directly from Appendix A, a self-paced course organized around the themes of the Advisory, and linked examples of policies and guidance from other jurisdictions.

The website also curates practical materials that boards and divisions can use or adapt. These include resources developed for this project and relevant publicly available materials from governments, education organizations, privacy authorities, and other reputable bodies. The collection can grow and change as needs develop and useful resources become available.

Readers should consult the companion website for the most current edition, updates, links, and supporting materials. It is available at: https://saskschoolboards.ca/ai.

3. The Context for This Policy Advisory

3.1 From generative AI to embedded infrastructure

AI did not enter education with ChatGPT. Schools already used automated systems for adaptive learning, search, plagiarism detection, learning analytics, accessibility, cybersecurity, and administration. ChatGPT's public release in late 2022 made generative AI widely accessible for producing text, images, audio, video, and code, exposing the need to validate these tools and establish age-appropriate, privacy-protective uses (UNESCO, 2023). Academic-integrity concerns remain important, but policy also needs to address AI's uses in planning, accessibility, translation, tutoring, communication, data analysis, and administration. Risks include inaccurate or biased content, obscured evidence of learning, and exposure of personal or confidential information.

AI features are increasingly embedded in word processors, search engines, email and collaboration platforms, learning systems, accessibility tools, and administrative software. They may arrive through routine product updates rather than separate procurement decisions. A division can therefore acquire new AI capabilities within software it already licenses, even when the board has never approved a stand-alone AI tool.

A blanket decision to permit or prohibit AI cannot address differences in technologies, purposes, users, and risks. Boards need durable principles, clear boundaries for prohibited and restricted uses, and administrative processes for evaluating tools, proposed uses, and material changes to approved systems.

3.2 Understanding the broader context

Several external conditions are shaping the decisions boards must make. They are included here as context rather than as matters any individual board can resolve.

An industry in rapid motion: A relatively small number of large technology companies provide much of AI's infrastructure, computing capacity, and foundation models. Capabilities, terms, pricing, and information practices can change quickly. Canada's national AI strategy identifies dependence on large global providers as an economic and sovereignty concern (Government of Canada, 2026). Divisions may therefore need to review material changes to a tool, its intended use, or its handling of information.

Expanding adoption and expectations: Canada's national AI strategy, launched in June 2026, includes measures to expand responsible adoption and AI literacy, including CanCode support for K–12 learners and educators (Government of Canada, 2026). These developments signal that AI literacy is increasingly being treated as part of broader educational, economic, and civic preparedness. They do not require every student or staff member to use AI, but they strengthen the case for helping people understand where it is useful, where it can fail, and when it should not be used.

An evidence base that is still developing: Much research on generative AI in education concerns short interventions, particular tools, or narrowly defined learning contexts. The OECD

notes that much of the available experimental research lacks sufficient statistical power or contextual detail and can be difficult to generalize, replicate, or scale (OECD, 2026). Boards should therefore avoid treating either early benefits or early harms as settled conclusions. A more durable approach is to require a clear purpose, appropriate safeguards, human oversight, and continued evaluation as evidence and technologies develop.

An evolving legal landscape: At the time this Advisory was prepared, Canada did not have a general federal AI statute. In June 2026, the federal government introduced Bill C-36, the Protecting Privacy and Consumer Data Act, and Bill C-34, the Safe Social Media Act. As of October 12, 2026, neither bill had received Royal Assent (Parliament of Canada, 2026a, 2026b). These proposals include measures concerning children's information, automated decision-making, online safety, and some AI chatbots (Bill C-34, 2026; Innovation, Science and Economic Development Canada, 2026b). Their provisions may change before enactment, and any new obligations will depend on the final legislation, its scope, regulations, and coming-into-force dates. These proposals do not replace the obligations Saskatchewan school divisions already have under provincial law. Their direction nevertheless reinforces the need for transparent, privacy-protective, and regularly reviewed local governance.

Environmental and resource costs: AI depends on physical infrastructure. The International Energy Agency projects that global electricity demand from data centres will more than double between 2024 and 2030, with AI as the most important driver of that growth (International Energy Agency, 2025). AI systems can also contribute to water use through data-centre cooling, electricity generation, and semiconductor manufacturing (UK Government Digital Sustainability Alliance, 2025). Because effects vary by system, provider, location, energy source, and scale of use, the Advisory addresses environmental responsibility through AI literacy, responsible use, and procurement. Participants at the Indigenous Educators Summit also identified water and energy demands as matters the education sector should consider (Couros, 2026).

3.3 The Saskatchewan context

The broader developments described above take on particular significance in Saskatchewan.

Different communities, capacities, and circumstances: Saskatchewan's 27 school divisions serve diverse urban, rural, northern, First Nations, and Métis communities, with differing enrolment, geography, connectivity, staffing, fiscal capacity, and access to specialized technical expertise. An approach manageable for a large urban division may be difficult to sustain in a smaller or geographically dispersed one. Survey respondents identified infrastructure and capacity pressures, particularly for small and rural divisions, as important operational concerns (Couros, 2026). This Advisory is intended to support locally appropriate implementation while identifying areas in which shared services or collective approaches may be more effective than parallel local efforts.

First Nations and Métis relationships and authority: Saskatchewan boards operate on Treaty territories and the Homeland of the Métis, often alongside First Nations education authorities that operate schools under their own authority. Where AI use involves Indigenous students, languages, knowledge, data, or community interests, general engagement with stakeholders may not be

sufficient. Divisions must identify and respect applicable relationships, rights, protocols, and Indigenous data governance requirements. Section 5 considers these responsibilities in greater detail.

Variable access to devices and connectivity: The guidance is intended to remain usable without one-to-one devices or continuous high-bandwidth access, including through shared devices, teacher-mediated activities, and low-tech approaches to learning about AI. It also reflects the provincial limits on student use of personal cellphones during instructional time, introduced for the 2024-25 school year (Government of Saskatchewan, 2024). Those limits do not remove AI from classrooms because AI capabilities may still be present in school-provided devices, approved platforms, and other educational technologies.

The Saskatchewan legal framework: A central statute for school-division AI governance is *The Local Authority Freedom of Information and Protection of Privacy Act* (LA FOIP). It governs school boards as local authorities and establishes requirements concerning access to records and the collection, use, disclosure, and protection of personal information. Boards exercise their broader educational authority under *The Education Act, 1995*. Depending on the proposed use, duties involving human rights, employment, copyright, child protection, records, and other areas of law may also apply. The Office of the Saskatchewan Information and Privacy Commissioner has published AI-related privacy resources, including materials on children's privacy and the use of AI (Office of the Saskatchewan Information and Privacy Commissioner, n.d.). Section 6.3 addresses the application of privacy requirements to AI governance.

An established digital citizenship foundation: The province's 2015 Digital Citizenship Education in Saskatchewan Schools guide and its Kindergarten to Grade 12 continuum provide a developmental approach to areas such as digital access, privacy protection, online rights and responsibilities, literacy, security, and well-being (Couros & Hildebrandt, 2015). This Advisory builds on that foundation while addressing the distinct capabilities and governance questions introduced by AI. The AI Literacy and Use Continuum, an illustrative governance reference available through the companion website, is informed by this earlier Saskatchewan work and by more recent AI literacy frameworks.

An uneven policy landscape: A December 2025 EdCan Network scan identified ministry-level K-12 AI guidance only in British Columbia, New Brunswick, and Quebec (EdCan Network, 2025). Within Saskatchewan, local activity and readiness have varied, while the SSBA consultation identified a need for shared guidance that boards could adapt to their own circumstances (Couros, 2026). This Advisory responds to that need while preserving the authority of each board to establish policy suited to its communities, responsibilities, and capacity.

Taken together, these conditions favour an approach based on durable principles, proportionate safeguards, local adaptability, and continuing review. Section 4 sets out the principles that provide that foundation.

4. Guiding Principles

This Advisory rests on nine guiding principles. The principles were synthesized from a review of more than 150 policy documents and frameworks on AI in education, including material from Canadian school divisions and provincial ministries and from national and international organizations. WAG members and partners then ranked the principles through the consultation survey, identifying human-centred education and student well-being as the leading concerns (Couros, 2026).

Each principle is expressed as a durable commitment intended to outlast particular tools and practices. The principles are not written as policy language that a board would adopt verbatim. Section 4.10 maps where each principle is developed throughout the Advisory, and Appendix A provides sample policy wording that boards may adapt.

4.1 Human-centred education

Education is a human endeavour built on relationships among students, teachers, families, communities, and the staff and leaders who support them. AI should support, not displace, those relationships or the professional judgement on which education depends. When considering a proposed use, boards and division leaders should ask how it may affect students, staff, families, communities, learning, relationships, and human judgement.

Responsibility for human-centred use is shared in ways appropriate to each role. Teachers apply the principle in planning, instruction, feedback, and assessment. Staff apply it when AI supports communication, records, scheduling, and other services. Students develop the judgement to use AI while strengthening their own understanding, independent thinking, and capacity to learn. AI should support this development without replacing students' own thinking. In practice, AI may inform decisions about people, but it should not replace accountable human judgement. Time saved through AI should be reinvested in learning, relationships, and work requiring human care and expertise.

The central test is whether a proposed use advances the purposes of public education and helps students grow as capable, discerning human beings. This principle comes first because every other principle is an application of it.

4.2 Student well-being and children's rights

Many general-purpose AI systems were not designed specifically for children or educational settings. When a division requires or facilitates their use, it must establish that the proposed use is suitable rather than assume that a commercially available product is appropriate. Decisions should account for age and developmental stage, mental health and well-being, privacy, participation, protection from harm, and the requirement that the best interests of the child be a primary consideration (United Nations, 1989; UN Committee on the Rights of the Child, 2021). UNESCO also recommends age limits for independent conversations with generative AI platforms and emphasizes age-appropriate safeguards (UNESCO, 2023).

Well-being includes both immediate safety and longer-term development. Boards should consider whether a proposed use supports independent thought and healthy relationships, whether its design

encourages prolonged or emotionally dependent use, and whether credible evidence supports its claimed educational benefit. Existing duties to respond when a student may be at risk continue to apply wherever a concern comes to light.

Section 6.1 develops the implications for AI literacy, development, and well-being. Section 6.5 establishes boundaries for harmful and higher-risk uses, including image-based abuse and student-facing conversational systems.

4.3 Privacy, security, and responsible data use

Personal information about students and staff requires the same protection whether it is held in an established information system or processed through a new AI-enabled service. LA FOIP applies to AI systems and embedded AI features just as it applies to other information systems. Boards remain responsible for personal information in their possession or under their control, including when a service provider processes that information on their behalf.

A proposed use should be assessed for what information it collects or generates, why that information is needed, where it is processed and stored, who can access it, how long it is retained, how it is deleted, and whether it may be used to train or improve a model. Privacy is addressed both when tools and uses are approved and when staff and students use them in daily practice.

Responsible data use also requires restraint. Divisions should collect and retain only what is necessary and consider whether the intended purpose can be achieved without entering personal information into an AI system. A use may be technically lawful and still be inappropriate because of the sensitivity of the information, the nature of the system, or the availability of a less intrusive approach. An approved tool is not necessarily approved for every use. Section 6.3 develops this principle in depth.

4.4 Equity, inclusion, and access

The benefits of AI should not depend on where a student attends school, and its risks and burdens should not fall most heavily on the schools and communities with the least capacity to manage them. Across Saskatchewan, connectivity, devices, staffing, specialist expertise, and budgets differ among urban, rural, remote, northern, and First Nations communities. Policy that assumes equal capacity may widen existing disparities.

Equity also includes accessibility and fairness. AI may improve access for students with disabilities and diverse learning needs, but systems may also perform unevenly across languages, dialects, cultures, disabilities, and demographic groups. A useful test is whether a proposed initiative could be implemented safely and meaningfully in the school facing the greatest implementation constraints. If it could not, the division should identify the supports or reasonable alternatives needed before proceeding. Section 6.6 expands on this principle.

4.5 Ethical and responsible use

Ethical and responsible use concerns both whether an AI system should be used and how people use it. UNESCO's Recommendation on the Ethics of Artificial Intelligence identifies proportionality and doing no harm as foundational principles (UNESCO, 2021).

First, AI should not be used to harm others. AI tools can be used to fabricate convincing images, voices, messages, and identities. Using AI to deceive, harass, humiliate, exploit, defraud, or impersonate a person in a harmful manner is not merely a technical misuse. The same applies to using AI to gain unauthorized access to systems, accounts, or information or to evade safeguards established by a division. Policy should prohibit such conduct and connect it to the division's existing safety, conduct, privacy, cybersecurity, and incident-response processes.

Second, AI use should be proportionate and appropriate to the purpose. The availability of a tool is not, by itself, a reason to use it. Expected benefits should justify the educational, human, financial, privacy, security, and environmental costs, including the energy and water demands discussed in Section 3.2. Deciding not to use AI may be an exercise of sound judgement when its costs or risks outweigh its benefits, when it does not serve the intended purpose, or when human authorship, relationship, or professional judgement is integral to the work.

The greater the potential harm to a student or employee, the stronger the expectations should be for evidence, transparency, privacy protection, meaningful human review, an effective way to question decisions and seek correction, and ongoing monitoring. This principle also applies to new uses not specifically anticipated in this Advisory. Section 6.4 applies proportionality to tool approval and procurement.

Some work also derives meaning from human authorship and relationship. A letter of condolence, a personal expression of thanks, or an acknowledgement of the territory on which a school stands should not be treated merely as a task to be completed efficiently. Section 6.5 establishes prohibited and restricted uses.

When approving or governing AI uses involving culturally significant language, stories, images, symbols, artefacts, or other cultural expressions, divisions should consider accuracy, context, attribution, applicable permissions or restrictions on access and reuse, and the risk of stereotyping or misrepresentation. For sensitive or consequential uses, administration should determine whether relevant cultural expertise or community engagement is needed before proceeding. This general responsibility complements, but does not replace, the distinct Indigenous rights and authority addressed in Section 4.7 and Section 5.

4.6 Teacher professional agency

AI policy should recognize teachers as professionals. Teachers should have a meaningful voice in decisions about tools and practices that shape classroom work. Within applicable curriculum, policy, and administrative requirements, professional agency includes the capacity to question, adapt, override, or decline AI-generated recommendations and to determine that an approved tool is unsuitable for a particular class, student group, task, or educational purpose. Sustained professional learning, adequate time, and access to support are necessary before staff are expected to use unfamiliar systems.

Policy should also protect teachers from AI being used to monitor or evaluate their work outside established professional processes. Systems that observe classroom activity or generate analytics

about instructional practice should not displace fair, contextual, and accountable approaches to professional evaluation.

Professional responsibility cannot be transferred to a tool. AI may assist with planning, feedback, or communication, but teachers remain responsible for the materials they use and the professional decisions they make. Assessment and evaluation of student learning must remain human professional acts. Sections 6.1 and 6.2 expand on professional learning, professional judgement, assessment, and evaluation.

4.7 Indigenous authority and self-determination

AI governance in Saskatchewan schools must respect Indigenous self-determination, data sovereignty, and the authority of First Nations and Métis rights holders over their knowledge, languages, cultural practices, and community-held information. Divisions must distinguish decisions that fall within their authority from decisions that require direction or authorization from the relevant First Nation, Métis government or community, or First Nations education authority.

Respecting Indigenous authority requires more than verifying the accuracy of AI-generated material or protecting individual privacy. An output may be accurate and comply with ordinary privacy requirements while still bypassing the permissions, relationships, protocols, collective interests, or authority governing the knowledge or information involved. Indigenous authority should therefore remain a standing governance consideration in decisions about teaching materials, data, procurement, analytics, and other AI applications.

Section 5 explains how this principle applies to language, ceremony, sacred knowledge, community-held information, decisions affecting students and staff, and continuing engagement.

4.8 Transparency and accountability

People affected by an AI use should receive information proportionate to its purpose and potential effect. Students should know when AI materially contributes to their learning or assessment and what role it plays. Families should be able to learn which approved student-facing tools are used, for what purposes, and what information those tools collect. Staff should know when their work, communications, or meetings are recorded, analyzed, or used to inform decisions about them. Boards should receive reporting on significant AI uses, risks, incidents, costs, and outcomes at least annually.

Transparency also applies to how students and staff use AI. Students should follow clear expectations for disclosing permitted AI assistance in submitted work. Staff should follow established expectations for reviewing and, where appropriate, disclosing substantive AI assistance in materials, communications, and decisions. These obligations differ by role and context, but both depend on honesty about the contribution AI has made.

When AI contributes to a consequential decision, an identifiable person must remain responsible for that decision. The person affected should be able to understand AI's role, identify an error, raise a concern, and seek meaningful human review or correction. Section 6.2 addresses disclosure in

assessment. Section 6.6 develops the requirements for public information, reporting, review, and accountable oversight.

4.9 Adaptability and continuous learning

AI technologies, products, and common practices will continue to change. Principles endure longer than the systems they govern, which is why policy should be written at the level of durable commitments and paired with a defined review process. Section 8 recommends an approach for periodically reviewing this Advisory.

AI literacy is a continuing responsibility for board members, division leaders, staff, and students rather than a one-time activity. Adaptable governance requires boards to revisit earlier assumptions, consider new evidence and experience, and recognize when a previously appropriate decision requires revision. It also requires clear routes through which staff, students, families, and communities can raise concerns or identify emerging needs.

Continuous learning includes humility. Some current judgements will need to change. Responsible governance means anticipating that possibility, maintaining the information needed to review decisions, and establishing processes through which policies, procedures, approvals, and safeguards can be improved.

4.10 How the guiding principles function throughout this document

The nine guiding principles recur throughout the Advisory. Sections 6.1 through 6.6 organize the six policy domains that boards must govern, and Section 6.7 applies those domains to division administration and business functions. Each domain draws on several principles at once.

The table below identifies where each principle is developed most fully and where important additional applications appear. The mapping is indicative rather than exhaustive. Human-centred education informs the entire Advisory, and the guidance concerning Indigenous authority in Section 5 applies across all domains.

Guiding principle	Developed chiefly in	Also appears in
4.1 Human-centred education	Sections 6.1 through 6.6; the organizing centre of Section 6	Indigenous authority and relationships (Section 5); administrative and business applications (6.7); the AI Literacy and Use Continuum on the companion website; Appendix B
4.2 Student well-being and children's rights	AI literacy and well-being (6.1); prohibited and restricted uses (6.5)	Required student use and reasonable alternatives (6.4); equity and student voice (6.6); the AI Literacy and Use Continuum on the companion website

Guiding principle	Developed chiefly in	Also appears in
4.3 Privacy, security, and responsible data use	Privacy, security, and data governance (6.3)	Tool approval and contracts (6.4); restricted uses and wearable devices (6.5); meetings, records, and administrative systems (6.7); Appendix B
4.4 Equity, inclusion, and access	Equity, accountability, and community engagement (6.6)	Saskatchewan context (Section 3); equitable AI literacy opportunities (6.1); required use and reasonable alternatives (6.4); the AI Literacy and Use Continuum on the companion website; Appendix B
4.5 Ethical and responsible use	Tool approval and procurement (6.4); prohibited and restricted uses (6.5)	Wider social and environmental context (3.2); AI literacy (6.1); academic integrity and assessment (6.2); administrative and business uses (6.7)
4.6 Teacher professional agency	AI literacy and professional learning (6.1); assessment and evaluation (6.2)	Surveillance and consequential decisions (6.5); human resources and workforce change (6.7)
4.7 Indigenous authority and self-determination	Indigenous authority, knowledge, and AI governance (Section 5)	Sections 6.1 through 6.6, with specific applications in data governance, procurement, and prohibited or restricted uses (6.3 through 6.5); Appendix B
4.8 Transparency and accountability	Equity, accountability, and community engagement (6.6); disclosure in assessment (6.2)	Privacy notices and records (6.3); documented approval and review (6.4); restricted uses (6.5); administrative and business functions (6.7); Appendix B
4.9 Adaptability and continuous learning	Keeping This Policy Advisory Current (Section 8)	Continuing review of AI systems and approvals (6.3 and 6.4); routes for questions and concerns (6.6); annual self-check and priority action planning (Appendix B)

TURNING PRINCIPLES INTO POLICY

A board does not need nine separate policies to honour nine principles. Most boards will address them most effectively through one concise board policy that states the board's commitments and expectations, supported by administrative procedures maintained by the division. The sample policy statements in Appendix A are designed to support that structure, and the companion website provides a guided policy generator and related templates.

5. Indigenous Authority, Knowledge, and AI Governance

Some uses of AI involve Indigenous knowledge, languages, cultural practices, or community-held information over which a school division does not have sole authority. These questions may arise when staff use AI to prepare learning materials, teach or translate Indigenous languages, generate text or images involving Indigenous cultural elements, or process records already held by the division. They may also arise when AI features are added to software the division already licenses.

The central governance question is not simply whether a proposed use concerns Indigenous Peoples. It is whether the use falls within the division's authority or requires direction, agreement, authorization, or consent from the relevant First Nation, Métis government or community, or First Nations education authority. Board policy should establish this boundary and provide a basis for administration to develop clear guidance for staff. This Advisory governs only decisions that a board or division is authorized to make. It does not direct how Indigenous communities use AI in their own language, cultural, or data work.

Indigenous Peoples and communities do not hold a single position on AI. Participants at the Indigenous Educators Summit described useful applications in planning, differentiation, and administrative work, particularly when time saved supports relationships with students and land-based learning. They also identified firmer boundaries involving language, ceremony, and sacred knowledge. Participants emphasized that First Nations and Métis contexts differ and that the June 4, 2026 gathering began a conversation that requires continued engagement (Couros, 2026). Board policy should therefore uphold Indigenous authority and decision-making rather than treat AI as either inherently appropriate or inherently inappropriate in Indigenous education. The Summit gathered First Nations and Métis perspectives, and this section does not claim to represent Inuit perspectives.

5.1 Direction on Indigenous language, ceremony, and sacred knowledge

The clearest board-level direction concerns uses that a division should not initiate independently. The WAG consultation survey favoured weaving Indigenous perspectives throughout the Advisory, while the Indigenous Educators Summit identified a clear boundary around uses of generative AI involving Indigenous language, ceremony, ceremonial protocol, and sacred knowledge (Couros, 2026). This guidance does not define authority on behalf of First Nations or Métis communities. It identifies circumstances in which a division should not assume that it holds the authority to proceed.

This approach is consistent with Article 31 of the United Nations Declaration on the Rights of Indigenous Peoples, which recognizes Indigenous Peoples' rights to "maintain, control, protect and develop their cultural heritage, traditional knowledge and traditional cultural expressions" (United Nations, 2007, p. 9).

BOARD-LEVEL DIRECTION

A division should not use AI to teach an Indigenous language or to generate, reproduce, or prepare material involving ceremony, ceremonial protocol, or sacred knowledge, including for lesson planning or resource preparation, unless the relevant First Nation, Métis government or community, or First Nations education authority has directed or authorized the use. This direction governs division decisions and does not limit decisions Indigenous communities make about their own knowledge, languages, or cultural work. If a division is uncertain whether a proposed use falls within this area, it should not proceed until it has sought direction from the appropriate rights holder.

This direction is not a blanket prohibition on every use of AI involving Indigenous Peoples or subject matter. Summit participants distinguished uses involving general planning, publicly documented history, and other information that may fall within division authority from uses involving language, ceremony, ceremonial protocol, or sacred knowledge. Uses within division authority still require professional judgement, reliable and appropriately sourced information, and verification of generated content (Couros, 2026). Section 6.1 addresses the professional learning needed to apply this distinction, while Section 6.5 explains how it relates to prohibited and restricted uses.

5.2 Accuracy does not establish authority

AI-generated Indigenous content raises an immediate concern about accuracy. Generative systems can fabricate facts, sources, identities, and language while presenting the result confidently. The risk is especially serious when the user lacks the knowledge needed to recognize an error. Summit participants warned that an educator who lacks confidence with Indigenous content may treat an AI system as an advisor rather than rely on authentic sources and appropriate relationships. They also described systems inventing Indigenous nations and attempts to obtain ceremonial names through AI rather than through the relationships and protocols by which those names are given (Couros, 2026).

Generated content should therefore be checked against authentic and vetted sources and, where appropriate, knowledge shared through recognized relationships. Verification is necessary, but it answers only whether the content is reliable. It does not establish whether the division had authority to request, generate, receive, or use it.

Authority is a separate question. Whether Indigenous knowledge may be shared can depend on its source, the relationships involved, the permission provided, and applicable community protocols. An AI system cannot obtain that permission or establish those relationships for a division. Langley School District similarly advises staff not to use AI to create Indigenous learning opportunities because AI cannot follow Indigenous protocols, seek permission, identify where teachings originated, or limit itself to authentic sources (Langley School District, n.d.).

Boards should therefore distinguish two questions: Is the content reliable, and did the division have authority to produce and use it? An output may be accurate while its production or use still bypasses the people, permissions, relationships, and protocols that govern the knowledge. Accuracy does not cure a lack of authority.

5.3 Community knowledge and data in division custody

School divisions may hold materials that raise questions beyond individual privacy. Archives and program records may include recordings of Elders and Knowledge Keepers, oral histories, locally developed language resources, cultural materials, and other community-contributed information. These records may engage collective interests and authority as well as the privacy rights of identifiable individuals.

The Local Authority Freedom of Information and Protection of Privacy Act governs personal information, but compliance with the Act does not by itself determine whether a division has authority to upload community-contributed material to an AI system, permit a vendor to retain it, summarize it automatically, or use it to train or improve a model. Transcription, summarization, and other AI processing may be embedded in platforms a division already uses. These activities can therefore occur without a separate procurement decision unless the division actively reviews and governs embedded features.

For First Nations data, the principles of OCAP® provide an important governance framework. The principles establish how First Nations data and information should be collected, protected, used, and shared. They address First Nations ownership and control of their information, access to information about themselves wherever it is held, and possession or stewardship through physical control of data (First Nations Information Governance Centre, n.d.). Each First Nation determines how the principles apply within its own context. OCAP® is specifically a First Nations framework and should not be treated as a generic framework for all Indigenous Peoples.

For Métis or Inuit information, a division should not substitute OCAP® for the requirements of the relevant rights holder. It should work with the relevant Métis government or community, Inuit rights-holding organization, or other appropriate body to identify the applicable governance requirements. This does not suggest that the Advisory has defined those requirements on their behalf.

Before using community-contributed information or knowledge with AI, a division should determine whose information is involved, where it came from, the purpose for which it was provided, what permission accompanied it, whether the proposed use differs from that purpose, who holds authority over subsequent uses, and who may benefit or bear risk. Possession does not establish authority. These questions of provenance, purpose, permission, authority, benefit, and risk should be addressed alongside privacy, security, and contractual requirements. Sections 6.3 and 6.4 expand on these responsibilities.

5.4 Decisions affecting students and staff

Care is also required when AI informs decisions or judgements about First Nations and Métis students or staff. Some Summit participants cautioned against descriptions of Indigenous students that emphasize perceived shortcomings and warned that AI-assisted assessment may give conventional Western assessment methods a new appearance of objectivity (Couros, 2026). AI does not remove limitations or biases in the underlying data, instruments, or practices. It may reproduce them while making the resulting judgement appear more neutral or authoritative.

This concern applies to student assessment, predictive analytics, program placement, discipline, applicant screening, employee assessment, and other consequential uses. A person responsible for

the decision must understand the evidence, consider its limitations and context, and remain accountable for the outcome. AI should not be permitted to turn an existing assumption, incomplete record, or measure focused on perceived shortcomings into an apparently objective conclusion.

AI-detection tools are one such example. Language and dialect can influence their results, and research has reached different conclusions about whether particular detectors disadvantage writers whose first language is not English. No detection score should be treated as proof about an individual student's work. Section 6.2 examines this evidence and recommends a fair process based on conversation, demonstrated understanding, and other available evidence. Section 6.5 restricts high-stakes analytics and fully automated decisions, while Section 6.7 applies related safeguards to human resources and other administrative functions.

5.5 Engagement and continuing dialogue

Engagement is most useful when it begins before a decision has already been made. Not every AI matter requires engagement with First Nations or Métis rights holders, but divisions should identify the decisions that do. They should also distinguish general engagement from circumstances in which direction, agreement, authorization, or consent is required. Uses involving Indigenous language, ceremony, ceremonial protocol, sacred knowledge, or community knowledge and data may fall within the latter category. Where direction or authorization is required, the division should identify the relevant nation or community and follow its protocols for the particular knowledge and proposed use. Engagement with an advisory group or one community should not be assumed to authorize use on behalf of other rights holders.

Other decisions may not require formal authorization but still warrant early engagement because of their potential effects. These may include analytics, AI-enabled recruitment or screening, employee assessment, and other decision-support systems that could materially affect First Nations or Métis students or staff. Early engagement allows potential effects, assumptions, safeguards, and alternatives to be considered before adoption. Depending on the matter and the authority involved, a proposed use may be supported, rejected, modified, or deferred. Each is a legitimate governance outcome.

Engagement and policy development should also account for differences in access and capacity across Saskatchewan. Participants from northern divisions described infrastructure, staffing, and community realities that differ from those often assumed in technology initiatives. WAG consultation participants similarly cautioned against treating any one provincial context as typical (Couros, 2026). Boards should consider whether a proposed use depends on connectivity, devices, staff capacity, or professional learning that are not equally available across schools and communities. Section 6.6 develops these questions of equity, capacity, accountability, and engagement.

Future engagement should build on the June 4, 2026 Summit by including students, additional First Nations education authorities, more northern and remote communities, and other rights holders whose interests may be affected. Participants also asked the SSBA to continue convening the conversation as AI technologies and division practices evolve. Sections 6.1 through 6.7 translate the considerations in this section into guidance concerning professional learning, assessment, data governance, procurement, restricted uses, equity, accountability, community engagement, and administrative functions. Section 8 recommends an approach for keeping the Advisory current.

6. Policy Domain Guidance

Sections 1 through 5 establish the foundations for understanding and governing artificial intelligence, and Appendix C provides plain-language definitions. Section 6 is the practical policy core of this Advisory. Sections 6.1 through 6.6 address six connected policy domains identified through the WAG consultation (Couros, 2026) and informed by Canadian and international policy, guidance, and research. The domains are organizing lenses, not silos: a single proposed use of AI may engage several domains at once. Cross-references identify those connections while avoiding unnecessary repetition.

Each subsection follows the same core structure:

- an "In Brief" summary;
- essential background for board readers; and
- a reference to the corresponding board discussion questions, sample policy statements, and drafting note in Appendix A.

The background explains why an issue matters; it does not mean that every background consideration must appear in board policy. Appendix A provides printable board discussion questions and sample policy statements for each subsection, together with drafting notes explaining important wording choices and policy distinctions. The sample statements are examples, not requirements, and provide a common starting point for local policy development. Boards may select, combine, and revise them to fit local circumstances. Local adaptation must remain consistent with applicable law and any direction or authorization required from relevant Indigenous rights holders. Appendix B is an optional planning tool for identifying gaps, distinguishing board decisions from administrative work, and recording priority actions, lead roles, target dates, and evidence of completion. The companion website provides supporting resources, templates, and practical tools.

Two tenets apply throughout Section 6:

- Consistent with Principle 4.1, a human-centred approach considers how decisions about AI affect people, relationships, well-being, learning, professional judgement, and public trust.
- Governance, administration, and professional practice remain distinct. Provincial curricula and statutory responsibilities establish the broader educational framework. Boards establish policy commitments, boundaries, and oversight expectations. Division administration implements them through procedures, system approval, guidance, and support. Teachers exercise professional judgement in learning and assessment within that framework.

Teaching and learning is addressed across the six domains rather than treated as a separate domain. **Across these domains, boards can apply a common test to any proposed AI use: does it support student learning or the conditions for learning, strengthen professional judgement, and sustain public trust.** For learning-related uses, AI should support the thinking, interaction, creativity, and reflection through which students develop understanding and capability. Administrative uses may support these purposes indirectly by improving services or freeing time for teaching and relationships. Each domain applies a different governance lens, while the Advisory remains focused on board policy rather than classroom pedagogy.

The six policy domains are:

- AI literacy, professional learning, and well-being;
- academic integrity and assessment;
- privacy, security, and data governance;
- AI tools, procurement, and intellectual property;
- prohibited and restricted uses; and
- equity, accountability, and community engagement.

Section 6.7 is not a seventh domain. It applies the six domains to the division's administrative and business functions, including human resources, finance, facilities, transportation, communications, and records.

6.1 AI literacy, professional learning, and well-being

IN BRIEF

AI literacy enables students, staff, division leaders, and board members to understand AI, evaluate its outputs, use approved tools responsibly, and judge when AI should not be used. Boards build this capacity by setting clear expectations, supporting sustained professional learning, providing equitable access to AI literacy learning and support, and requiring age-appropriate safeguards that protect learning, relationships, well-being, and professional judgement.

Why AI literacy comes first

AI policy depends on people being able to interpret and apply it. Staff need enough understanding to protect information, select approved tools, explain expectations, and recognize when AI should not be used. Students need a shared language for discussing acceptable use and evaluating outputs. Board members and division leaders need sufficient literacy to make informed policy, procurement, and oversight decisions. AI literacy is therefore a governance capacity as well as a learning priority. AI literacy should develop alongside critical thinking, judgement, creativity, collaboration, and the ability to plan, monitor, and reflect on one's own learning.

What AI literacy includes

AI literacy is more than knowing how to operate a tool. It includes understanding, in general terms, how AI systems produce outputs and why those outputs can be inaccurate or biased; knowing how to use approved tools for a defined purpose; checking AI-produced claims against reliable sources; protecting one's own and others' privacy; and deciding whether a particular use supports the learning or work at hand. For board policy, AI literacy is a capacity goal rather than a prescribed course of study. Its application should be role-appropriate and consistent with provincial curricula, professional responsibilities, approved-system requirements, and division procedures.

Evidence about AI and learning

Research on how generative AI affects student learning remains limited and ongoing. A 2026 review from Stanford University's SCALE Initiative screened more than 800 papers in its K–12 AI research repository but identified only 20 with causal evidence strong enough to inform its main findings. The review cautions that the strongest studies still examine a narrow range of tools, subjects, and settings and that many conclusions remain preliminary (Fesler et al., 2026).

Recent randomized studies illustrate how outcomes can vary with the way AI is designed and used. In one peer-reviewed field experiment, nearly 1,000 students in grades 9 through 11 at a large high school in Turkey were assigned to one of three conditions during mathematics practice: no AI; a standard GPT-4 interface designed to resemble ChatGPT, called GPT Base; or a GPT-4 tutor with teacher-informed safeguards that provided hints and limited direct answers, called GPT Tutor. Both AI groups performed better than the control group on the assisted practice problems. On a subsequent unassisted assessment, however, the GPT Base group scored 17 per cent lower than the control group, while the GPT Tutor group's performance was statistically indistinguishable from the control group (Bastani et al., 2025).

In a separate World Bank working paper, researchers evaluated a six-week after-school program for first-year senior secondary students in nine public schools in Benin City, Nigeria. Students assigned to the program used Microsoft Copilot for curriculum-aligned English activities twice weekly, with teachers introducing topics, guiding students' use, and supporting reflection. Participating students outperformed the control group on an end-of-program assessment covering English, AI knowledge, and digital skills (De Simone et al., 2025).

These studies were relatively short and conducted in particular settings, so their results should not be assumed to apply across subjects or contexts. For boards, the central implication is that access alone does not establish educational benefit. **Improved AI-assisted output is not, by itself, evidence of improved learning.** Reviews of learning-related pilots, products, and benefit claims should consider evidence of what students understand, retain, and can apply. Division administration should evaluate proposed systems and implementation conditions, while decisions about instructional purpose and classroom use remain matters of professional practice within provincial curricula and division policy.

Frameworks for AI literacy

Established frameworks can help boards understand the capacity required for responsible AI governance and help division administration plan role-appropriate learning and support. Student learning remains governed by provincially authorized curricula and teachers' professional judgement. UNESCO's student framework identifies four competency areas: a human-centred mindset, ethics of AI, AI techniques and applications, and AI system design. It describes three progression levels: Understand, Apply, and Create (UNESCO, 2024a). UNESCO's teacher framework identifies five areas: a human-centred mindset, ethics of AI, AI foundations and applications, AI pedagogy, and AI for professional learning. It uses three progression levels: Acquire, Deepen, and Create (UNESCO, 2024b).

In Saskatchewan, the University of Saskatchewan has adapted the Digital Education Council's AI literacy framework into separate resources for university students, faculty, and staff. The student framework is organized into five dimensions: understanding AI and data; critical thinking and judgement; ethical and responsible use; human-centricity, relationality, and creativity; and domain expertise and career readiness. It also describes three levels of progression within each dimension (University of Saskatchewan, 2025). These resources were designed for a post-secondary institutional context and include expectations related to coursework, research, professional practice, and career readiness. They should not be treated as a K–12 developmental framework. They are nevertheless a useful Saskatchewan reference for shared institutional language, role-specific

learning, progression, relationality, and the kinds of AI-related knowledge and judgement students may increasingly encounter in further education and professional settings.

Together, the UNESCO and University of Saskatchewan resources illustrate the knowledge, skills, and judgement that AI literacy can involve in different educational contexts. The companion website provides an illustrative, grade-banded AI Literacy and Use Continuum as a governance reference for boards. It draws on these and other frameworks to help board members consider developmental differences, ask informed questions, and assess whether division policy, safeguards, supports, and review processes are adequate. The continuum does not prescribe curriculum or classroom practice. Learning about AI is distinct from direct use of an AI system: progression through the grade bands does not authorize access or imply that older students should use AI tools more often or more independently. Student-facing use remains subject to division approval and the safeguards in Sections 6.4 and 6.5.

Professional learning and capacity

Relevant, meaningful, and ongoing professional learning is a condition of responsible AI implementation. It should help staff connect AI use to curriculum, assessment, privacy, approved tools, student well-being, and professional judgement. Professional learning should be role-specific, available to teachers, administrators, support staff, division leaders, and board members, and updated as tools and risks change. Board member professional learning should include practice in applying the common test set out in the introduction to Section 6 to proposed AI uses and assessing the evidence for the expected benefits. It should be supported by time, guidance, and opportunities to practice and learn with colleagues, both before unfamiliar systems are introduced and as approved uses evolve. It should also prepare staff to apply Section 5, Indigenous Authority, Knowledge, and AI Governance, when AI use involves Indigenous content, knowledge, language, ceremony, or cultural practice.

Student rights, access, and well-being

Student rights and well-being are important considerations when setting boundaries for AI use. The Convention on the Rights of the Child and General Comment No. 25 confirm that children's best interests, privacy, participation, and protection from harm apply in digital environments (United Nations, 1989; UN Committee on the Rights of the Child, 2021). For boards, this means that decisions about access should consider age, developmental readiness, educational purpose, supervision, privacy, and evidence of benefit.

Equitable access is also an AI literacy issue. As AI increasingly shapes learning, further education, employment, public services, and civic life, students who lack opportunities to understand and critically evaluate it may be less prepared to participate safely and effectively in those settings. Equity does not require every student to use the same tool. It requires every student to have age-appropriate opportunities to learn about AI, evaluate it critically, and receive support suited to local circumstances. Where connectivity, staffing, devices, or local capacity are limited, the division should provide appropriate alternatives so that access to AI literacy learning does not depend on a student's school or community.

Conversational and companion-style systems require particular caution because they can simulate care, authority, or companionship. For generative AI platforms generally, UNESCO recommends a

minimum age of 13 for independent conversations, together with age verification and parental or guardian monitoring (UNESCO, 2023). A division may establish stricter requirements and should also consider product terms, privacy obligations, developmental readiness, and educational purpose. Conversational or companion-style systems should not replace trusted adults, peer relationships, or established student-support services. If an AI interaction reveals a safety or well-being concern, staff should follow the division's existing response and reporting procedures. Section 6.5 establishes the approval, monitoring, and safety requirements for student-facing conversational and companion-style systems.

Educator well-being

Implementation of AI systems also affects educator well-being and professional judgement. AI may save time in some tasks, but poorly planned adoption can add training, checking, and administrative work or introduce forms of monitoring that undermine trust. Boards should involve staff in implementation, communicate the purpose and limits of approved uses, and ensure that AI does not replace professional evaluation or decision-making. Where AI creates efficiencies, those benefits should support teaching, relationships, and well-being rather than creating new workload expectations. Board reporting should account for verification and correction demands, new or less visible work, how workload is redistributed, and effects on professional agency, alongside any claimed time savings.

Board discussion and policy development: See Appendix A.1 for the board discussion questions, sample policy statements, and drafting note associated with this subsection.

6.2 Academic integrity and assessment

IN BRIEF

Work submitted for assessment is ordinarily expected to provide evidence of a student's learning. When generative AI has materially shaped that work, the finished product may not clearly show what the student contributed or what the student knows, understands, or can do independently. The board should establish division-wide policy principles for AI use in learning and assessment, including disclosure, fair academic-integrity processes, approved systems, and accountable human judgement. Administration should implement those principles through procedures, guidance, professional learning, and shared resources, while teachers retain professional judgement within provincially authorized curricula and division policy. The evidence does not justify procuring AI-text detection tools to determine misconduct, and a detector score should not be treated as a finding or used as the sole basis for an allegation or decision.

The assessment and integrity challenge

Saskatchewan guidance identifies three purposes of assessment: assessment for learning, which informs teaching and supports improvement; assessment as learning, which involves students in monitoring their own learning; and assessment of learning, which provides evidence of achievement. Together, these purposes contribute to an overall picture of a student's learning (Saskatchewan Ministry of Education, 2022).

Generative AI does not change these purposes, but it can make completed work more difficult to interpret. A polished essay, presentation, solution, or other product may reflect independent student ability, permitted AI assistance, or substantial reliance on AI. Assessment design should help

teachers distinguish what a student can produce with AI from what the student knows, understands, and can do. This creates two related governance concerns: students need clear information about division-wide requirements and about how AI may be used or disclosed for a particular task, and decisions about possible misconduct must be fair and based on relevant evidence.

Responsibility is shared but distinct. The province establishes authorized curricula and provincial requirements. The board administers and supervises the school division and establishes policy. Administration develops procedures, approves systems, and provides guidance and support. Teachers exercise professional judgement in planning learning, designing assessment, and evaluating student achievement within that framework (*The Education Act, 1995*, ss. 85, 86 and 231; Saskatchewan Ministry of Education, 2022).

Division-wide clarity

Students should receive clear, accessible, and timely information about both division-wide expectations for AI use and any school-, course-, or task-specific requirements. The information should make clear which expectations apply generally and which depend on the particular learning or assessment context.

Board policy can require division administration to establish common language and guidance for communicating these expectations. Schools can make division-wide expectations accessible through handbooks, responsible-use materials, orientation, and communication with families. Teachers communicate how the expectations apply to particular learning and assessment tasks, including whether and how AI may be used and whether disclosure or citation is expected. This delivery chain can reduce unnecessary inconsistency while preserving teachers' professional judgement.

The AI Assessment Scale, developed in higher education, is one example of a five-level framework for communicating different degrees and purposes of permitted AI use (Perkins et al., 2024). Administration may adapt an existing framework or develop simpler language suitable for Saskatchewan K–12 schools. Selection and implementation of such a framework are administrative matters, not prescribed elements of board policy.

Division guidance may also distinguish disclosure from citation. Disclosure explains whether and how AI contributed to the work. Citation identifies a source of information, language, media, or ideas. Whether either is required for a particular task depends on its purpose and the applicable professional expectations.

Reliable and fair decisions

Saskatchewan assessment guidance emphasizes evidence gathered over time and attention to validity, reliability, and fairness (Saskatchewan Ministry of Education, 2022). For board governance, the relevant principle is that an academic-integrity decision should rest on relevant evidence and professional judgement rather than on an automated inference. The evidence appropriate to a particular task or learning cycle remains a matter of professional practice. Administration is responsible for establishing procedures through which concerns are documented, considered, and resolved.

AI-text detectors estimate whether writing displays patterns associated with generated text. They do not establish who produced the work, whether AI use was permitted, what the student understands, or which parts were completed independently.

Published studies have reached different conclusions about detector performance. Weber-Wulff et al. (2023) found substantial variation among 14 tools, including false positives, false negatives, and weaker performance after editing or paraphrasing. Liang et al. (2023) found frequent false positives for a sample of TOEFL essays written by people whose first language was not English, raising an important equity concern. However, its exploratory design and mismatched comparison groups limit how broadly the result can be applied. Jiang et al. (2024), using purpose-built detectors and a much larger set of prompt-matched GRE essays, achieved high accuracy and found no evidence of disadvantage to writers whose first language was not English in that setting.

These studies tested different detectors, datasets, and conditions. They are not direct replications, and success in a controlled study does not establish that a commercial detector will perform reliably across varied K–12 work. Taken together, the research does not justify procuring AI-text detection tools to determine whether a student engaged in misconduct. An incorrect allegation can create inequitable consequences, damage trust, and harm a student's relationship with teachers and school. A detector score should not be treated as a finding of misconduct or used as the sole basis for an allegation or decision.

Board policy should require academic-integrity processes that protect confidentiality, consider relevant evidence, and provide the student with a meaningful opportunity to respond. Canadian school-board guidance similarly recommends conversation and consideration of evidence from products, observations, and conversations rather than reliance on a detection result (Educational Collaborative Network of Ontario [ECNO] & Ontario Association of School Business Officials [OASBO], 2025). Administrative procedures should define responsibilities, documentation, escalation, and any review or appeal process.

Those procedures should distinguish among misunderstanding a communicated requirement, failing to make a required disclosure, and deliberately submitting AI-generated content as the student's own contrary to the requirements for the task. Board policy may require responses to be educative and proportionate to the student's age, intent, understanding, previous instruction, and the effect of the conduct.

Accountable human judgement

Teachers' professional judgement remains central to assessment, evaluation, and communication of student achievement (Saskatchewan Ministry of Education, 2022). AI may assist with preparing assessment materials, drafting feedback, or organizing information relevant to evaluation when an approved system is used under division requirements. Before AI-assisted material is communicated as feedback or used in evaluation, the teacher must review it against the student's work and other relevant evidence, correct or reject it where necessary, and make the final judgement. The teacher remains accountable for the feedback or evaluation communicated to the student.

Students should be informed when AI materially contributes to individualized feedback, including whether the feedback was reviewed and communicated by a teacher or delivered directly by an approved instructional system. An approved instructional system may provide immediate formative feedback under conditions established by administration, and students should be able to bring that feedback to a teacher for review. Student work must be assessed and evaluated by a teacher, and no summative evaluation or final decision about student achievement should be made solely by an automated system. This is consistent with Canadian school-board guidance stating that AI-generated feedback may supplement but not replace professional judgement and that student work is assessed and evaluated by teachers (Halton District School Board, 2025).

Board discussion and policy development: See Appendix A.2 for the board discussion questions, sample policy statements, and drafting note associated with this subsection.

6.3 Privacy, security, and data governance

IN BRIEF

School divisions remain responsible for personal information when AI systems or AI-enabled features are used. *The Local Authority Freedom of Information and Protection of Privacy Act* (LA FOIP) applies to these uses as it does to other information systems. The board should establish division-wide policy requirements for approved use, privacy review, protection of personal information, provider agreements, incident response, and oversight. Administration should implement those requirements through documented reviews, agreements, inventories, safeguards, procedures, training, and continuing oversight. Where information has collective First Nations or Métis dimensions or includes Indigenous knowledge or community-contributed material, the division must also identify and respect the applicable Indigenous data-governance authority.

LA FOIP applies to AI

School divisions hold sensitive information about students, families, and staff. This can include educational records, disability and accommodation information, health and counselling information, family circumstances, attendance and behavioural records, employee information, and information created through the use of digital systems.

LA FOIP defines a board of education or conseil scolaire as a local authority. Its requirements apply whenever the division's use of an AI system or AI-enabled feature involves personal information. Using an external system, an outside provider, or a feature embedded in an existing platform does not reduce the division's responsibility.

LA FOIP regulates the collection, use, and disclosure of personal information and requires local authorities to maintain appropriate administrative, technical, and physical safeguards. Under section 23.1 of LA FOIP, those safeguards must protect the integrity, accuracy, and confidentiality of personal information and address reasonably anticipated threats, loss, and unauthorized access, use, disclosure, or modification. If personal information is used or disclosed without authority, section 28.1 of LA FOIP requires the division to take all reasonable steps to notify the affected individual where it is "reasonable in the circumstances to believe that the incident creates a real risk of significant harm" (LA FOIP, s. 28.1).

Board policy need not prescribe every safeguard or response step. It should require the division's handling of personal information through AI systems and AI-enabled features to comply with LA

FOIP, assign accountability, and require administration to maintain appropriate controls, procedures, training, incident response, and reporting.

Personal information and approved use

Everyday AI use can involve personal information even when a user does not intend to disclose it. Names and direct identifiers are not the only concern. A prompt, uploaded file, conversation history, output, or combination of details may identify a person. An AI system may also create an inference, classification, summary, or prediction about an identifiable individual.

Personal information should not be entered into an AI system that has not been assessed and approved for that purpose. Approval should apply to both the system and the proposed use. An AI system approved for staff brainstorming with non-confidential information is not thereby approved to process student work, counselling notes, psycho-educational reports, employee records, or information used to develop or revise an Inclusion and Intervention Plan (IIP).

Where an approved use involves personal information, only the information reasonably necessary for that purpose should be used. Staff should use division-managed accounts, settings, and workflows where required and should not move personal information into consumer accounts or unapproved systems. The Alberta School Boards Association recommends centralized vetting and approved-tool lists. Ontario guidance from ECNO and OASBO similarly directs staff to consult approved-tool lists, limit the information submitted, and exclude identifiable or confidential information from prompts (Alberta School Boards Association, 2024; ECNO & OASBO, 2025).

Risk-based privacy review

Before approving a new or materially changed use of an AI system that handles personal information, administration should use the division's established privacy-review process to document the purpose and legal authority, the necessity and proportionality of the information involved, relevant information flows and safeguards, applicable notice or consent requirements, and the possible effects on affected individuals. Particular care is required where sensitive information, children, monitoring, profiling, or consequential decisions are involved.

Canadian privacy regulators recommend privacy and related impact assessments as methods for identifying and mitigating risks from generative AI, with particular attention to children and other vulnerable groups (Office of the Privacy Commissioner of Canada and provincial and territorial privacy authorities, 2023). The Saskatchewan Information and Privacy Commissioner's AI resource collection brings together guidance from Saskatchewan and other Canadian privacy regulators, including material on children's privacy and AI use (Office of the Saskatchewan Information and Privacy Commissioner, n.d.). The depth of review should reflect the sensitivity and scale of the information, the people affected, and the potential consequences. Divisions can use existing screening and privacy-impact-assessment processes rather than creating a separate process for every AI use.

Information management service providers

Using an outside provider does not transfer the division's obligations. Under section 23.2 of LA FOIP, a local authority must enter into a written agreement before disclosing personal information to an information management service provider. The agreement must govern access, use, disclosure, storage, archiving, modification, and destruction; protect the personal information; and require

compliance with LA FOIP and its regulations. Section 8.2 of the Regulations also requires a description of the service, security and safeguarding obligations, and provisions for destruction where applicable.

For an AI-enabled service, the agreement and system settings should specify how the existing statutory requirements apply to each AI-enabled function. In particular, they should specify whether personal information may be used for model training or service improvement and the purposes and limits of any permitted use. Broader procurement requirements, including vendor selection, renewal, portability, transition, and exit planning, are addressed in Section 6.4.

A material change to an AI system's information practices, model, integrations, subcontractors, or approved purpose should trigger renewed privacy and security review.

Security, incidents, and continuing oversight

Privacy and security are related but distinct. Privacy concerns whether personal information is collected, used, disclosed, and retained with proper authority and limits. Security concerns how information and systems are protected from loss, disruption, unauthorized access, manipulation, or attack. Both are required for responsible AI use.

AI can add accounts, integrations, permissions, and data flows to a division's technology environment. It can also increase and strengthen phishing, impersonation, social engineering, and other cyber threats. The Canadian Centre for Cyber Security notes that threat actors may use AI to disguise themselves or imitate someone known to the intended victim (Canadian Centre for Cyber Security, 2026b).

The division should incorporate approved AI systems into its existing access-control, cybersecurity, records-management, business-continuity, and enterprise-risk practices. Cybersecurity and incident-response processes should also address AI-enabled threats. Administration should define operational responsibilities, escalation, investigation, containment, recovery, notification, documentation, and review.

Approval is not final. Administration should reassess or, where necessary, suspend an approved AI use when an incident, credible complaint, unexpected system behaviour, change in purpose, or material provider change could alter the basis on which it was approved. The board should receive reports on significant risks, incidents, and corrective actions through its established risk and incident-reporting processes. Administration remains responsible for technical investigation, operational response, and decisions about immediate controls.

Indigenous data governance

Some information in division custody has collective or community dimensions that are not resolved by compliance with individual privacy law. Examples may include recordings of Elders or Knowledge Keepers, oral histories, language resources, cultural materials, community-contributed records, or information about lands, histories, or collective experience. Possession of this material does not by itself give a division authority to submit it to an AI system, permit a provider to retain it, or use it to train or improve a model.

For First Nations data and information, the principles of OCAP® provide an important governance framework. OCAP® addresses ownership, control, access, and possession and establishes how First

Nations data and information should be collected, protected, used, and shared. Each First Nation determines how the principles apply within its own context. OCAP® is specifically a First Nations framework and should not be treated as a generic framework for all Indigenous Peoples (First Nations Information Governance Centre, n.d.).

For Métis information, a division should work with the relevant Métis government or community to identify the applicable authority and governance requirements. More generally, before using Indigenous knowledge or community-contributed information with AI, the division should determine whose information is involved, why it was provided, what permission accompanied it, whether the proposed use differs from that purpose, and who has authority over subsequent use. These requirements supplement the privacy, security, and contractual obligations that otherwise apply. The broader governance principles are addressed in Section 5.3.

> **Board discussion and policy development:** See Appendix A.3 for the board discussion questions, sample policy statements, and drafting note associated with this subsection.

6.4 AI tools, procurement, and intellectual property

> **IN BRIEF**
>
> AI systems and AI-enabled features may enter a division through purchases, pilots, free accounts, staff requests, or changes to software already in use. Board policy should require administration to manage them through established technology and procurement processes from initial proposal to retirement. Administration's review should address the purpose and expected value of the system, the division's ability to implement and sustain it, the provider relationship, intellectual property, and eventual exit. Privacy and information-handling requirements are addressed in Section 6.3, while prohibited and restricted uses are addressed in Section 6.5.

Managing acquisition and the service lifecycle

AI procurement is broader than purchasing new software. Free tools, trials, classroom pilots, and AI features added to software the division already licenses can also create legal, technical, financial, and support obligations. Board policy should therefore require one documented pathway for considering proposed systems and uses through implementation, renewal, and retirement.

Administration should integrate the pathway with existing purchasing and technology processes, assign responsibilities, and maintain a record of each decision. The record should identify the system or service version, purpose, accountable administrator or unit, authorized users or groups, conditions of approval, and the date or circumstances that require reassessment. It should also incorporate any additional review required under division policy rather than create separate records for each policy domain.

Because some uses may initially be unknown or inadvertent, discovery of an unrecorded or unapproved use should trigger a proportionate administrative response. Administration should assess the use, suspend or discontinue it where required, address any resulting risk or harm, and update guidance or controls where unclear expectations or a system gap contributed. The response should distinguish good-faith or inadvertent use from deliberate non-compliance or harmful conduct.

As part of this technology and procurement review, administration should answer three questions. What educational or operational purpose would the system serve, and what evidence supports the expected benefit? Can the division implement and sustain it, considering accessibility, compatibility, support, workload, and full cost? What legal, policy, and contractual requirements must be addressed? These questions synthesize considerations found across Canadian school-system guidance (Alberta School Boards Association, 2024; British Columbia Ministry of Education and Child Care, n.d.; ECNO & OASBO, 2026).

Administration should determine the depth of review and expertise required. More extensive review is appropriate when a proposal involves students, an essential function, a substantial financial commitment, or significant long-term dependency.

Where a pilot is appropriate, administration should define its purpose, duration, accountable administrator or unit, safeguards, and evaluation criteria. Administration must ensure that the pilot meets the privacy, security, accessibility, restricted-use, contractual, and other division requirements that apply to the proposed use. At the end of the pilot, administration should record its decision: proceed with the proposed use, revise it, discontinue it, or conduct further review.

AI acquisitions remain subject to the division's established purchasing requirements. Depending on the nature and total value of the contract, this may include requirements for an open and competitive procurement process. School divisions may reduce duplication by sharing assessments, model contract language, accessibility findings, or joint purchasing arrangements with one another. Each participating division remains responsible for deciding whether the system and proposed use meet its own requirements.

Provider relationships, material change, and exit

The division's legal obligations when a provider handles personal information are addressed in Section 6.3. As part of procurement, administration must also address the broader service relationship. Agreements should define the service and support to be provided, accessibility and account-administration responsibilities, who owns or may use prompts, uploaded files, outputs, and other material handled through the service, information available for oversight, notice of material changes, pricing and renewal, portability, transition, termination, and any obligations that continue after the service ends.

Staff should not accept provider terms on behalf of the division outside the approved process. Nor should labels such as "enterprise," "education," or "secure" be treated as evidence that the agreement, settings, and account controls meet division requirements.

Before entering an agreement, administration should plan for foreseeable service changes and for the eventual end of the relationship. The agreement should require notice when the provider changes material features, models, terms, data-handling practices, or prices. It should also establish how the division can retrieve its information, transfer the records and system settings needed to continue operations with another service, and verify deletion when the relationship ends. The Canadian Centre for Cyber Security's guidance addresses data rights, migration, retention, and destruction in cloud-service agreements (Canadian Centre for Cyber Security, 2024).

Replacing a service can become difficult or costly when information is stored in a format that other systems cannot readily use, the service is connected to many division systems, staff workflows

depend on it, or the provider charges substantial migration fees. Administration should consider these forms of dependency before adoption and again before renewal.

When a provider changes a model, feature, term, data-handling practice, or price, administration should determine whether continued use remains appropriate or requires further review. Before renewing an agreement, administration should review continuing need, performance, cost, dependency, material changes, and available alternatives.

Copyright, ownership, attribution, and publication

AI raises two related questions: what material may be submitted to a system, and how AI-assisted output may be used. Uploading copyrighted or licensed material to an external AI system may reproduce the material or disclose it to a provider in a way that the original classroom use, purchase, or licence did not authorize. Any submission must therefore be consistent with applicable law, licences, agreements, provider terms, and the approved purpose.

Section 29 of the federal *Copyright Act* includes education among the permitted purposes for fair dealing, but an educational purpose does not make every dealing fair. The Council of Ministers of Education, Canada provides guidance concerning copying protected material for students (Council of Ministers of Education, Canada, n.d.). Permission to use material in one educational setting does not necessarily authorize its submission to an external AI provider. Staff should follow established copyright and fair-dealing procedures and refer significant or uncertain uses for review.

AI-assisted output may reproduce protected expression, misstate attribution, or include material whose source and permitted use are unclear. Official material intended for publication or broad distribution should be reviewed for accuracy, provenance, attribution, permissions, and possible infringement.

Canadian questions concerning AI training, human authorship, ownership, and responsibility for infringing output continue to develop. A federal consultation found disagreement on these issues and reported that Canadian courts had not yet examined allegedly infringing AI-generated content (Innovation, Science and Economic Development Canada, 2025). Provider terms may allocate contractual rights, but they do not settle every question under Canadian copyright law. Board policy should not declare that every AI output is automatically owned by the user or the division.

Ownership and permission require particular care when student-created or staff-created work is submitted to an AI system. Under the federal *Copyright Act*, the author is generally the first copyright owner. A student therefore generally retains copyright in original work the student creates; submitting the work for instruction or assessment does not by itself transfer ownership to the division. Division procedures should identify when student work may be uploaded, what students and families will be told, whether permission is required, and whether the provider may retain or reuse the work.

The position for staff-created work depends on the circumstances. Under subsection 13(3) of the *Copyright Act*, the employer is generally the first copyright owner when an employee creates a work in the course of employment, unless an agreement provides otherwise. Collective agreements, employment terms, assignments, licences, and the circumstances in which the work was created may affect ownership and permitted use. Section 14.1 also provides authors with moral rights that cannot be assigned, although they may be waived. Division policy should therefore avoid assuming that

every teacher-created work belongs either to the teacher or to the division. Procedures should identify who may authorize submission, reuse, licensing, or publication in the circumstances.

Division procedures should also distinguish disclosure of AI assistance from citation or attribution. Staff and administrators remain responsible for copyright, licensing, permissions, confidentiality, attribution, and publication review.

Ordinary intellectual-property rules do not determine authority over Indigenous knowledge, language, stories, cultural expressions, or community-contributed material. Division approval cannot replace the direction or authorization of the relevant rights-holding First Nation, Métis government or community, or First Nations education authority. Provider agreements must not authorize training, service improvement, product development, or another use of such material without that authority.

Required student use and reasonable alternatives

Learning about AI does not necessarily require students to create an account or submit work to a commercial system. Before requiring students to use a particular AI system directly, the division should approve that system and the proposed use for the relevant age or grade group. The approval process should address age and account requirements, privacy, consent, accessibility, accommodation, supervision, support, and reasonable alternatives. These governance conditions do not determine how a teacher designs a particular learning activity.

Required use must comply with *The Saskatchewan Human Rights Code, 2018*, including the duty to accommodate disability-related needs to the point of undue hardship (Saskatchewan Human Rights Commission, 2023). Where consent has not been provided or the system cannot be used safely or accessibly, a reasonable alternative may be required. The alternative should allow the student to meet applicable course or program requirements without academic disadvantage arising from the inability to use the system.

Halton District School Board provides one Canadian example. Its guidelines require signed parent or guardian consent for specified classroom uses of approved generative AI systems and require an alternate assignment when consent is not provided (Halton District School Board, 2025). This is a local board approach, not a general legal rule for every Canadian school division or every AI-enabled feature.

> **Board discussion and policy development:** See Appendix A.4 for the board discussion questions, sample policy statements, and drafting note associated with this subsection.

6.5 Prohibited and restricted uses

> **IN BRIEF**
>
> This section distinguishes prohibited uses, which should not occur within the scope of division policy, from restricted uses, which may proceed only for an approved purpose and with defined safeguards. It addresses Indigenous authority, intimate or sexualized depictions, harmful impersonation, protected material, automated decisions and learning analytics, biometric and behavioural technologies, AI-enabled wearables, conversational systems, and AI agents.

Background

Clear boundaries allow boards to adapt this guidance to local circumstances. Prohibited-use provisions should identify the conduct that is not allowed, the scope of the prohibition, and any narrow exception or authorization. Restricted-use provisions should identify the approved purpose, the person authorized to approve it, the required safeguards, who remains accountable, and when the approval must be reviewed.

In their principles for generative AI, Canadian privacy authorities recommend no-go zones for uses involving unlawful handling of personal information, discriminatory profiling or categorization, or purposes known or likely to cause significant harm or threaten fundamental rights (Office of the Privacy Commissioner of Canada and provincial and territorial privacy authorities, 2023).

Whether a use is prohibited or restricted depends on its purpose and effect, not only on the product. For example, a division might approve a biometric system solely to verify identity. It should not use facial expressions, eye movements, voice, or similar signals to infer a person's honesty or emotional state and then use that inference as a basis for judgement or decision-making.

A product may also perform more than one function. The same platform might converse with users, analyze student information, and take actions within connected systems. Each function should be assessed separately. Where functions overlap, the safeguards should reflect the function that presents the greatest risk and the consequences if it fails.

Prohibited uses and uses requiring specific authority

Some uses should not occur within the scope of division policy. Others may occur only through approved practices or with specific authority from a rights holder. The guidance below identifies prohibited uses, restricted uses, and any narrow authorization or exception that applies.

Indigenous authority

Section 5 explains that First Nations and Métis knowledge is governed by rights, relationships, protocols, and collective authority. A division should not use AI to teach an Indigenous language or to generate, reproduce, or prepare material involving ceremony, ceremonial protocol, or sacred knowledge unless the relevant rights-holding First Nation, Métis government or community, or First Nations education authority has directed or authorized that use. Permission from an individual user, vendor, or person who entered material into a system cannot substitute for that authority.

Intimate or sexualized depictions

AI can be used to create or alter intimate or sexualized depictions of a real person, including depictions for which no authentic source image exists. The Privacy Act (Saskatchewan) addresses the non-consensual distribution of intimate images and expressly includes altered visual recordings within its definition. Whether the Act applies to a particular AI-generated depiction depends on the facts and the statutory requirements.

Division policy should prohibit the creation, solicitation, or distribution of AI-generated or AI-altered intimate or sexualized images, audio, or video when the conduct involves a student, uses division systems, occurs through a school-related activity, involves an employee acting in a professional role, or affects safety, employment, or the learning environment. The division should respond with the seriousness applied to other forms of image-based abuse, support those harmed,

and obtain legal advice where required. Privacy authorities have identified non-consensual intimate imagery and other harmful AI-generated content depicting real people as a serious concern, particularly where children and other vulnerable groups are affected (Office of the Privacy Commissioner of Canada and international data protection authorities, 2026).

Harmful impersonation and deception

AI should not be used to impersonate an identifiable person in a manner intended or reasonably likely to deceive, harass, exploit, defraud, or cause harm. Examples include fabricating a message that appears to come from another person, cloning a voice to obtain information or money, or creating false media as supposed evidence of conduct. Satire, simulation, or instructional role-play should be clearly identified and should not be reasonably likely to harm or mislead.

Uploading and using protected material

Uploading material to an AI system may reproduce or disclose it to the provider and, depending on the account, settings, and agreement, may allow it to be retained or used to improve or train a model. Copyrighted, licensed, student-created, staff-created, and community-contributed material should therefore be submitted only through approved practices consistent with Sections 5, 6.3, and 6.4. AI should not be used knowingly to evade access or licence restrictions, remove or falsify attribution, or reproduce or distribute material contrary to law, licence, agreement, or division policy. Broader questions of fair dealing, ownership, attribution, provider terms, student and staff rights, and publication review are addressed in Section 6.4.

Unauthorized access and interference

AI should not be used to test or bypass security controls without authorization, obtain access to accounts or information, disrupt services, alter permissions, or interfere with division systems. These activities should be addressed through the division's existing cybersecurity, conduct, and incident-response processes.

Restricted systems and consequential uses

Restricted status attaches to a use, not simply to a product. Administrative procedures for restricted-use approval should require documentation of:

- the intended purpose and why a less intrusive approach is not sufficient;
- who will be affected and what benefits and harms are reasonably foreseeable;
- the evidence supporting the system's reliability and its known limitations;
- what information will be collected, generated, or inferred; where it will be stored or processed; who may access it; whether it may be used for model training; and how retention, backups, and verified deletion will be managed;
- what will be communicated to affected students, staff, or families;
- who will monitor the use and who can review, correct, or override an output; and
- the accountable administrator and the date or circumstances for renewed review.

The depth of review should be proportionate to the consequences of failure. Approval should be reconsidered when the purpose, users, system capabilities, information practices, or risks materially change. The privacy, security, contracting, and vendor-management requirements in Sections 6.3 and 6.4 apply to every restricted use. Section 6.1 establishes the professional-learning and capacity expectations that support responsible implementation.

Automated decisions and AI-assisted learning analytics

No fully automated system should make a final decision that significantly affects a person, including decisions about grades, discipline, placement, graduation, programs, services, accommodations, student support, or employment.

Requiring a person to approve a system's recommendation does not by itself provide meaningful oversight. The responsible person must have enough information, time, authority, and understanding to question the output, consider relevant circumstances, and reach a different conclusion. Affected people should have access to an understandable explanation and an appropriate process for correction or review.

Learning management systems and other educational platforms may include learning analytics that analyze attendance, participation, submissions, achievement, engagement, behaviour, or service use to generate patterns, classifications, alerts, predictions, or recommended interventions. Not all learning analytics use AI, but those using automated prediction, classification, or inference should be treated as AI-assisted analytics for this Advisory.

Canadian privacy regulators have cautioned that AI predictions about performance may affect a child's classroom experience and pre-emptively shape the child's educational path (Federal, Provincial and Territorial Privacy Commissioners and Ombuds with Responsibility for Privacy Oversight [FPT Privacy Regulators], 2025). A label such as "engagement," "risk," or "likelihood of success" should not be treated as a direct measurement of a student's effort, capacity, motivation, or future performance.

Analytics concerning identifiable students should therefore be restricted. Approval should require evidence that the analysis is suitable for its purpose and has been reviewed for reliability and bias. An AI-generated score, classification, alert, or recommendation should not be the sole basis for intervention, discipline, placement, access to services, or another consequential decision. Students and families should have an appropriate way to correct inaccurate information and raise concerns about how an analysis has been used.

Biometric identification, behavioural inference, e-proctoring, and surveillance

Biometric and behavioural technologies may be used for identity verification, online test supervision, attendance, engagement monitoring, safety, or student-support analytics. Canadian privacy regulators have identified online or remote proctoring systems, often called e-proctoring, that use facial recognition for authentication or monitor eye movements and keystroke patterns as educational-technology risks (FPT Privacy Regulators, 2025).

Biometric identification or verification should be restricted and require evidence of necessity, proportionality, accuracy, security, and compliance with privacy law. E-proctoring should be subject to the same review where it uses biometrics, behavioural monitoring, or recording. The availability of a feature does not establish that it is necessary or proportionate for a particular assessment.

Systems that claim to infer emotion, attention, honesty, distress, or intent from biometric or behavioural signals should not be used as a basis for assessment, discipline, employment decisions, or routine monitoring. This does not prohibit teachers from exercising professional judgement or divisions from considering observable participation information through approved learning analytics. It prohibits treating an automated inference about a person's internal state as a reliable

basis for judgement or decision-making. Such inferences may be uncertain, difficult to contest, and especially consequential where students or employees have limited power or choice.

The European Union's AI Act is not Saskatchewan legislation, but its prohibition on certain AI systems used to infer emotions from biometric data in educational institutions, except for specified medical or safety purposes, provides a useful comparative reference (European Union, 2024). Divisions should also prohibit covert or continuous AI-enabled surveillance beyond what existing law and division policy expressly authorize for a defined purpose.

AI-enabled wearable devices

AI-enabled wearables include glasses, watches, earbuds, pins, and other accessories that may combine cameras, microphones, location or biometric sensors with recording, transcription, translation, identification, content generation, or real-time prompting functions. Their use may be less apparent to others than the use of a phone or computer. Canadian privacy guidance notes that wearable devices can collect information about both users and people nearby and may record others without their knowledge (Office of the Privacy Commissioner of Canada, 2024).

Policy should address what a wearable is being used to do rather than its brand or appearance. Except for functions specifically approved as part of an accommodation or other authorized use, a wearable should not be used to record, stream, transcribe, identify, analyze, or transmit information about another person or to obtain assistance during an assessment or activity where that assistance is not permitted.

Wearables may support vision, hearing, communication, mobility, translation, or other medical and learning needs. School divisions have a duty to accommodate students with disabilities to the point of undue hardship (Saskatchewan Human Rights Commission, 2023). Policy should protect access to effective accommodation without guaranteeing approval of a particular device. The division should assess the student's needs, the device's functions, and risks to other people's privacy and rights through its accommodation process. Where those risks cannot be adequately addressed, the division may restrict or decline the device while providing an effective alternative that meets the student's accommodation needs. An approved use should identify which functions may operate, what information may be captured or transmitted, and how the privacy of others will be protected. Any restriction during an assessment or sensitive activity should maintain effective accommodation.

Conversational and companion-style AI

Section 6.1 addresses the literacy, developmental, relational, and well-being considerations associated with conversational AI. The policy question here is whether, and under what safeguards, a conversational system may be made available to students, staff, families, or the public.

Administrative approval of a student-facing instructional chatbot should be limited to defined educational purposes under age-appropriate conditions. Approval should address educational purpose, adult supervision, privacy, the sources and limitations of responses, whether interactions are retained or reviewed, and what is communicated to students and families. The system should identify itself as AI and should not be presented as authoritative or error-free.

Divisions may also deploy chatbots on public websites, portals, learning platforms, or internal service sites. A chatbot that provides general information without collecting personal information may present relatively little risk. A chatbot should be treated as a higher-risk restricted use if it

accepts identifiable questions, accesses records, operates inside an authenticated system, recommends services, creates or changes records, or handles matters involving attendance, discipline, accommodations, mental health, safety, complaints, or employment. Higher-risk approval should address authorized information sources, accuracy, privacy, security, data retention, logging, human escalation, correction of inaccurate responses, and accountable ownership. Every service chatbot should provide a visible way to contact a person.

Companion-style systems require stronger restrictions because they are designed or used to sustain an ongoing relational interaction. UNICEF identifies risks including emotional dependency, manipulation, or exploitation. It also reports dangerous responses in crisis-level conversations from some chatbots used for mental health support or therapy (UNICEF, 2025). Companion-style AI should not be provided or represented as a substitute for counselling, crisis response, peer relationships, or trusted adult support.

Students may disclose abuse, self-harm, threats, acute distress, or other serious concerns even when a chatbot was not designed for that purpose. A division should not assume that automated detection will recognize every serious disclosure. If a system uses safety alerts, procedures should define alert conditions, recipients, response periods, after-hours arrangements, record access, and what students and families are told. The system should not suggest that monitoring guarantees an immediate response, and the interface should provide clear directions for obtaining human or emergency help.

AI agents

AI agents require additional safeguards because they may access information, use tools, communicate, or make changes without a person completing each step. Approval should define the agent's purpose, permitted access, allowed actions, required human authorizations, activity records, monitoring, and procedures for stopping the agent.

The greater the possible consequence of an action, the less autonomy the agent should receive. Interim practical advice on agentic AI security similarly emphasizes limiting permissions, maintaining human oversight, monitoring activity, and retaining a reliable means of stopping an agent (National Cyber Security Centre, 2026). Section 6.7 applies these safeguards to administrative agents used across division operations.

Responding to misuse and harm

Clear boundaries should help divisions teach expectations rather than foster an enforcement culture based primarily on surveillance and punishment. Responses to student misuse should generally be educative, proportionate, and attentive to age, intent, understanding, impact, and previous instruction.

Conduct that harms another person, deliberately compromises protected material, or interferes with division systems requires a stronger response. Image-based abuse, harassment, exploitation, harmful impersonation, copyright or licensing misuse, unauthorized access, and deliberate interference with division systems should be addressed through the applicable safety, conduct, privacy, intellectual-property, cybersecurity, and incident-response processes. The division should support those harmed, preserve procedural fairness, meet applicable reporting obligations, and involve police or other authorities where required.

Incidents should also prompt review of whether unclear expectations, inadequate supervision, weak controls, or unmet learning needs contributed to the harm. Accountability for individual conduct does not remove the division's responsibility to improve its policies and safeguards.

> **Board discussion and policy development:** See Appendix A.5 for the board discussion questions, sample policy statements, and drafting note associated with this subsection.

6.6 Equity, accountability, and community engagement

> **IN BRIEF**
>
> This domain concerns whether the division's approach to AI remains fair, understandable, and accountable in practice. The board should receive enough information to identify division-wide patterns in access, benefit, cost, workload, risk, and burden; determine whether policy commitments are being met; and provide direction where change or additional resources are required. Administration should assign responsibility, maintain accessible public information, provide routes for questions and concerns, and engage affected communities in proportion to the significance of a proposed decision. The board should apply comparable expectations to its own use of AI.

Background

The preceding domains establish expectations for literacy and professional learning, assessment, privacy and data governance, system approval and procurement, and prohibited or restricted uses. This domain does not repeat those reviews. It asks whether their combined implementation remains equitable, visible, and accountable across the division as technology and local circumstances change.

Equity, accountability, transparency, and engagement reinforce one another. Unequal effects may remain hidden when affected people lack understandable information or a practical way to raise concerns. Engagement has limited value when responsibility for responding is unclear. The board cannot assess whether its policy is working without proportionate reporting. UNICEF similarly treats oversight, non-discrimination, transparency, accountability, and inclusion as connected requirements for child-centred AI (UNICEF, 2025).

Equity in access, benefit, and burden

Equity does not require every school, student, or staff member to use the same AI systems. It requires attention to whether people can benefit on fair terms, whether barriers and support needs are being addressed, and whether costs, risks, or additional workloads fall disproportionately on particular schools, groups, or communities.

AI systems are not interchangeable. Models, products, and account tiers may differ in capability, accessibility, language support, safeguards, usage limits, and cost. A division-wide initiative may therefore remain inequitable if students, staff, or schools receive materially different systems or service levels. Administration should examine what people can actually access and use, not only whether some form of AI is available.

The board should be able to determine whether a division-wide initiative can be implemented safely and meaningfully in schools with different levels of connectivity, equipment, staffing, specialist expertise, and administrative capacity. If implementation depends on resources that are not consistently available, administration should identify the necessary support, an appropriate alternative, or a narrower scope before proceeding division-wide.

AI may improve access through functions such as captioning, translation, text-to-speech, speech-to-text, or other forms of assistance. It may also create barriers when a system is inaccessible or performs unevenly across languages, dialects, disabilities, cultures, or demographic groups. Access to a system is not meaningful when a person cannot use it effectively, safely, or without giving up protections available to others. The United Nations Committee on the Rights of the Child connects meaningful digital access with non-discrimination and children's ability to participate in matters affecting them (UN Committee on the Rights of the Child, 2021).

Local capacity also affects cumulative cost and workload. Across several AI initiatives, licensing, infrastructure, professional learning, review, support, accessibility, records management, and transition requirements may redirect resources among schools or programs. Section 6.4 addresses these matters when a particular system or use is reviewed. Section 6.6 asks whether their combined effect is sustainable and whether the resulting distribution of resources requires board attention. Implementation plans should identify the funding required to sustain approved uses and any gap requiring additional funding. Where adequate resources are unavailable, boards should consider collaboration or advocacy for additional support, or adjust the scope and timing of implementation.

National reporting continues to identify connectivity gaps affecting rural, remote, northern, and First Nations communities (Canadian Radio-television and Telecommunications Commission, 2025). The WAG consultation also identified differences in local capacity and the need for approaches that Saskatchewan divisions of different sizes can reasonably implement (Couros, 2026). Some barriers can be addressed by a division. Others may require shared evaluation, interdivisional collaboration, coordinated procurement, common services, or advocacy for provincial support.

Administrative accountability and reporting to the board

The board establishes policy commitments and the information it requires for oversight. Division administration is responsible for procedures, training, documentation, operational decisions, and corrective action. Because AI governance may involve curriculum and instruction, student services, privacy, information technology, cybersecurity, human resources, finance, procurement, communications, records management, and legal advice, administration should assign responsibility for coordination at an organizational level with sufficient authority to obtain relevant information, involve appropriate expertise, pause a concerning use, and escalate significant matters. This need not create a new position, committee, or departmental structure.

At least annually, administration should provide the board, through existing reporting structures where appropriate, with concise information that enables it to:

- understand significant changes in approved and restricted AI uses;
- review serious incidents, recurring concerns, and corrective actions;
- assess evidence of benefits, limitations, costs, and unintended effects;
- identify material differences in access, accessibility, capacity, workload, risk, or impact among schools or affected groups; and
- decide whether changes in policy direction or resource allocation are required.

Annual reporting is a minimum cycle. A serious incident, material change in risk, or significant change in how an AI system is used should be reported sooner through established processes. Reports should protect personal information and avoid details that could compromise security.

Board members do not require an operational inventory of every minor use. Reporting should help the board identify patterns, test whether policy commitments are being met, allocate resources, and determine whether its direction should change. Administration remains responsible for technical investigation and day-to-day operational decisions.

Reporting should distinguish adoption (who uses AI and for what), implementation quality (whether the approved use is supported and safeguards are applied), and impact (its effects on learning, well-being, equity, workload, and organizational effectiveness). Adoption counts alone do not establish benefit. Claims of impact should identify the evidence, its limitations, and what remains uncertain.

Public transparency and routes for concern

Using information maintained through the division's approved-use process, students, staff, families, and community members should be able to find understandable information about how AI is used in the division. Public information should identify relevant policies, approved student-facing systems and purposes, the general categories of information those systems handle, important expectations for users, and a contact for questions or concerns.

The amount of explanation should reflect the significance of the use. A system that interacts directly with students, handles sensitive information, affects access to a service, or contributes to a consequential decision requires more explanation than a low-risk administrative aid. Transparency does not require disclosure of personal information, privileged or confidential records, protected contractual information, or information that could compromise security.

Transparency also applies when staff or administrators use AI. People should be informed when AI materially contributes to content, a recommendation, or a decision that affects them and knowledge of that assistance is relevant to trust, accountability, or their ability to question the result. Administrative procedures may distinguish these uses from routine spelling, formatting, transcription, or similar assistance that does not materially shape the content or judgement communicated.

AI may generate substantial public interest or disagreement. The board does not need to resolve every wider debate about AI, but the division should be able to explain why a local use is being considered, what evidence and uncertainty inform the decision, what benefits and risks have been identified, what safeguards or alternatives are available, and which matters fall within or beyond the division's authority. UNESCO's Recommendation on the Ethics of Artificial Intelligence supports proportionate transparency, accountability, human oversight, inclusive participation, and attention to human rights (UNESCO, 2021).

There should be a visible route for asking questions, reporting concerns, requesting correction of inaccurate personal information, or seeking review through the applicable division process. Canadian privacy regulators have emphasized the accountability of education authorities for technologies used in schools, clear information for students and families, and meaningful, age-appropriate participation by children and youth (FPT Privacy Regulators, 2025).

Community engagement and student voice

Engagement should occur early enough to influence a decision when a proposed use is broadly applied, difficult to reverse, likely to affect rights or access, or expected to alter learning conditions,

professional roles, or public trust. Depending on the issue, participants may include students, staff, families, school community councils, accessibility advocates, and affected cultural or community groups. Participants should be able to express support, concern, uncertainty, or opposition and identify effects that may not have been apparent to decision-makers. The board or administration, according to its authority, remains responsible for the decision and for explaining its basis.

Students should be included when a decision materially affects their learning, information, opportunities, or well-being. Participation should be appropriate to age and context, supported by understandable information, and conducted without unnecessary collection of personal information. The United Nations Committee on the Rights of the Child states that children's views should be sought and given due weight in decisions concerning the digital environment (UN Committee on the Rights of the Child, 2021).

Meaningful engagement includes reporting back about what was heard and how the input affected the decision. It does not require public consultation for every technology or operational choice. The form and extent of engagement should reflect the likely effect of the decision and the degree of community interest or concern.

General community engagement does not replace the distinct authority of First Nations and Métis governments, communities, and education authorities. Where rights-holder direction or authorization is required, the division should identify and follow that requirement rather than treat the matter as ordinary consultation.

The board's own use of AI

The board is also an AI user. Board members and people preparing board materials should protect confidential, privileged, personal, and in-camera information; use approved systems for board business; review AI-assisted content; and remain accountable for decisions and communications issued in the board's name.

The board should establish when material AI assistance in reports, presentations, correspondence, or other governance materials must be disclosed. Board member learning and board self-evaluation should also address AI governance so that the board can provide informed oversight and model the accountability it expects from administration.

> **Board discussion and policy development:** See Appendix A.6 for the board discussion questions, sample policy statements, and drafting note associated with this subsection.

6.7 Applying the domains: division administration and business functions

> **IN BRIEF**
>
> The six policy domains apply to every division function, not only to teaching and learning. This section applies them to human resources, finance, purchasing, meetings, records, communications, facilities, transportation, technology operations, and other administrative work. The board should establish the expectation that these uses are identified, approved for a defined purpose, proportionate to their consequences, subject to accountable human judgement, and monitored for their effects. Administration should translate that direction into operational procedures and controls appropriate to each function.

Applying the six domains to division operations

This section does not create a seventh policy domain or repeat the reviews required in Sections 6.1 through 6.6. It identifies how those requirements operate when AI is used to manage people, public funds, records, communications, services, physical systems, and essential division operations.

AI may enter administrative work through a stand-alone product, an employee account, or a feature added to software the division already licenses. All such uses, including embedded features not marketed primarily as AI, should be included in the division's approved-use record rather than tracked through a separate administrative inventory.

The consequences of administrative uses vary. An approved system may provide relatively routine assistance with drafting, summarizing, scheduling, translation, research, or data organization. Greater scrutiny is required when a system recommends, ranks, predicts, communicates, or takes action in a way that may affect employment, payment, access to services, resource allocation, safety, legal rights, or another significant interest. Approval must attach to the proposed purpose, not simply to the product. Approval for routine drafting does not authorize applicant screening, employee evaluation, payment approval, or another consequential use.

Human review must reflect the consequences of the use. For routine assistance, checking the accuracy, appropriateness, and confidentiality of the resulting work may be sufficient. Where an output informs a consequential decision, the responsible person must understand its purpose and material limitations, consider relevant information and circumstances, be able to question or reject the output, and remain accountable for the judgement or action taken.

Human resources and workforce change

AI may assist human resources staff with drafting job descriptions, developing interview questions, scheduling, organizing application materials, summarizing approved notes, or identifying information for further review. These uses do not transfer responsibility for a fair, accessible, and defensible employment process to the system.

Additional safeguards are required when AI screens or ranks applicants, interprets responses, recommends candidates, evaluates employee performance, or predicts qualities such as suitability, engagement, or future success. The process should use job-related criteria, be examined for barriers and bias, provide appropriate notice and routes for review, and preserve meaningful human judgement. Federal public-service guidance provides a useful, although not binding, reference: hiring managers remain accountable, must validate AI outputs, address bias and barriers, explain decisions, and tell candidates when and how AI will be used (Public Service Commission of Canada, 2026).

Candidates may also use AI to prepare application materials, respond to screening questions, or obtain assistance during an assessment or interview. Divisions should communicate what is permitted at each stage and use assessment methods that provide credible evidence of the qualifications being assessed. A suspicion of prohibited AI use should be considered through a fair process and should not be determined solely by an AI-detection score. The Public Service Commission advises that using an AI detector to assess candidate conduct is itself an automated assessment and requires corresponding scrutiny (Public Service Commission of Canada, 2026).

Video-conferencing and asynchronous interview platforms may include recording, transcription, translation, note generation, remote supervision, or analysis of voice, video, and screen activity. Candidates should receive advance notice of the functions being used, the information captured and retained, who may access it, and whether an output will affect their assessment. Before an AI-generated transcript, note, or summary is relied upon in an employment decision, its material content should be checked against the available recording and interviewers' notes. Federal guidance identifies facial-expression and tone-of-voice analysis as examples of assessment data that may not be transparent to candidates (Public Service Commission of Canada, 2026). A division should not use those or similar non-verbal signals to infer suitability, honesty, emotion, or attentiveness. The recording, confidentiality, and records requirements described below also apply to interviews and employment discussions.

Workplace implementation may affect duties, workload, required skills, professional judgement, staffing arrangements, monitoring, or evaluation. Before introducing a use that may materially affect employees, administration should assess its purpose and expected benefit, the reliability and limits of the system, possible bias or barriers, privacy and accessibility, effects on employees, and less intrusive alternatives. Administration should involve affected staff and bargaining agents as appropriate and identify the professional learning, role clarification, or other support required.

Finance, purchasing, and fraud prevention

AI may assist with routine financial and purchasing work such as drafting descriptions, organizing information, comparing documented requirements, summarizing records, matching transactions, or preparing preliminary analyses. Its output may contain incorrect calculations, invented facts, omitted conditions, or unsupported assumptions. A responsible person should verify source information, calculations, and material conclusions before the work informs a budget, payment, contract, forecast, or recommendation.

AI also increases the realism and scale of impersonation and social-engineering attempts. An email, voice message, video call, invoice, or other document may appear to come from a trusted leader, employee, supplier, or financial institution while containing a fraudulent request. The Canadian Centre for Cyber Security advises that AI can make phishing, personalized lures, cloned voices, and deepfake impersonation more convincing and recommends verifying sensitive or unusual requests through an independent communication channel (Canadian Centre for Cyber Security, 2026b).

Financial controls should therefore require independent verification for unusual or material requests, especially changes to supplier or employee banking information, urgent transfers, new payment instructions, or departures from established processes. AI assistance must not weaken delegated authority, segregation of duties, documentation, auditability, or independent review. An email, text message, collaboration-platform message, displayed telephone number, voice recording, video call, or AI-generated document should not by itself be treated as proof of identity or authority, even when the sender name, email address, telephone number, voice, image, or existing message thread appears familiar.

Where AI assists with the evaluation of bids, proposals, or suppliers, the division should protect confidential material, use the criteria established for the purchasing process, document the human evaluation, and preserve the authority of the person or body responsible for the decision.

Procurement of the AI system itself remains subject to the approval, contracting, intellectual-property, and service-lifecycle requirements established elsewhere in this Advisory.

Meetings, records, and access to information

AI recording, transcription, and summarization may improve accessibility and reduce administrative work. They may also create recordings, transcripts, prompts, summaries, and related information containing personal, confidential, privileged, or sensitive material. Before a tool is used, the division should identify the purpose, authority, information to be collected, access, storage, retention, and relationship to the official record.

Participants should receive advance notice that a meeting will be recorded, transcribed, summarized, or analyzed and should be told why the tool is being used. Where recording is not necessary and a participant objects, the division should consider a reasonable alternative. Where recording is necessary for an authorized purpose, that purpose and the available means of participation should be explained. Federal public-sector guidance similarly recommends advance notice, limiting recording to what is necessary, and respecting a participant's express objection where possible (Innovation, Science and Economic Development Canada, 2026a).

Staff should not introduce unapproved personal meeting bots into division meetings. AI recording, transcription, summarization, translation, or analysis should not be used during an in-camera board session, whether the session is online, hybrid, or in person, unless the use has been expressly authorized and the privacy, security, confidentiality, and records-management requirements have been addressed. The same restriction should apply to confidential personnel discussions, privileged legal discussions, and other protected meetings.

AI-generated notes may omit qualifications, combine speakers' views, or present uncertain statements as settled conclusions. A responsible person should check a summary against the meeting evidence before it is distributed or relied upon. Approved minutes remain the official record of board and committee decisions. Recordings, transcripts, prompts, summaries, and outputs may themselves be division records and must be managed accordingly. AI may assist staff in locating or organizing records for an access request, but a responsible person must determine whether the search is adequate and what information may lawfully be disclosed or withheld.

Official communications, translation, and public-facing services

AI may assist staff in preparing notices, reports, correspondence, presentations, website content, social-media posts, and other official material. An accountable staff member should review material released or relied upon in the division's name for accuracy, context, accessibility, confidentiality, copyright, and suitability for the intended audience.

AI assistance should be disclosed when it materially contributes to an official communication, report, recommendation, or record and knowledge of that assistance is relevant to trust, accountability, or the ability of an affected person to question the result. Procedures may distinguish these uses from routine spelling, formatting, transcription, or similar assistance that does not materially shape the content or judgement communicated.

Machine translation may improve access for families and community members, but it can alter meaning or omit important context. Communications involving safety, consent, student services,

legal rights, discipline, employment, or other significant consequences should be checked by a qualified person or through another reliable verification process. An accessible human-supported option should remain available where required.

A public-facing chatbot or virtual assistant should be clearly identified as an automated system, operate within an approved and tested scope, draw from current authoritative information, and provide an accessible route to a person. The division remains responsible for information delivered through a system it deploys. In *Moffatt v. Air Canada*, the British Columbia Civil Resolution Tribunal held an organization responsible for inaccurate information provided through its website chatbot. The decision is not binding on Saskatchewan school divisions, but it illustrates the accountability associated with deploying an automated public-information service (*Moffatt v. Air Canada*, 2024).

If a public-facing or internal service chatbot accepts open-ended communication, approval should address whether and how messages suggesting abuse, self-harm, threats, acute distress, or another urgent concern will reach an appropriate person. The system must not imply that it is continuously monitored or suitable for emergencies unless that is true. Where immediate monitoring is not provided, the limitation and appropriate emergency or human contact options should be clearly communicated. The restricted-use requirements for conversational systems remain in Section 6.5.

Operational analytics and physical systems

AI may support enrolment projections, transportation planning, scheduling, energy management, predictive maintenance, inventory management, facilities planning, and other operational analysis. These uses may assist in identifying patterns and comparing alternatives, but the analysis remains dependent on the relevance and quality of its information and assumptions. Historical data may reflect past inequities, changing community conditions, or service patterns that should not simply be reproduced.

An analysis that may significantly affect school boundaries, transportation, accessibility, staffing, facilities, or the allocation of resources should identify its sources, assumptions, limitations, and relevant alternatives. A person with appropriate responsibility must remain accountable for the analysis and recommendation. Where the matter requires a board decision, board members should receive enough information to understand how AI contributed and to question the recommendation without being expected to examine technical detail.

AI connected to building systems, access controls, cameras, sensors, vehicles, or other physical infrastructure requires security and safety review proportionate to the possible consequences. It should not independently alter a safety-critical system without defined authority, testing, monitoring, and a reliable means of human intervention. Requirements concerning biometrics, behavioural inference, surveillance, and consequential automated decisions continue to apply in these operational settings.

Administrative automation and AI agents

AI agents may retrieve or move information, update records, schedule activities, send messages, interact with other systems, initiate transactions, or complete a sequence of actions with limited human involvement. Their risk depends on both their purpose and the systems, information, and actions to which they have access.

Approval should define the agent's purpose, permitted information and systems, allowed actions, required human authorizations, activity records, monitoring, accountable owner, review cycle, and reliable means of suspension. Agents should receive only the access required for the approved task. They should not independently make consequential decisions about people, authorize payments, change access permissions, delete official records, or communicate externally on behalf of the division. Higher-risk agents should be tested in a controlled environment before deployment.

Cybersecurity, continuity, and evidence of value

Administrative and educational functions increasingly depend on connected division and provider-hosted systems. A service failure, provider change, or loss of access may interrupt payroll, student records, transportation, communications, facilities, finance, or other essential operations. A data breach may also expose information. The division should address both the security of its AI-enabled systems and the use of AI by people attempting to impersonate trusted individuals or gain unauthorized access.

Administration should identify essential functions, important provider dependencies, acceptable interruption periods, recovery priorities, alternate communication channels, tested backups, and workable fallback processes. Plans should assign responsibilities and address how the division will operate when a system or provider is unavailable. The Canadian Centre for Cyber Security recommends practiced incident-response and disaster-recovery plans, assigned roles, clear communications, and recovery arrangements established before an incident (Canadian Centre for Cyber Security, 2026a).

Claims that an administrative use will save time, reduce costs, improve service, or support better decisions should be evaluated against evidence. Evaluation should account for the time and resources required to verify outputs, correct errors, respond to concerns, train staff, administer the system, maintain parallel processes, and discontinue or replace the service. Material findings about uses that may significantly affect people, public funds, services, or essential operations should be included in the division's established reporting to the board.

> **Board discussion and policy development:** See Appendix A.7 for the board discussion questions, sample policy statements, and drafting note associated with this subsection.

7. Learning from Other Jurisdictions

Boards can learn from how other education authorities have translated broad principles into policy, administrative guidance, review processes, and practical tools. These resources can help identify recurring governance questions, reveal different approaches, and reduce unnecessary duplication of effort. They should be used as comparators rather than adopted unchanged. Each was developed within a particular legal, curricular, governance, and community context, and those conditions may differ substantially from Saskatchewan's.

The detailed annotated examples and links to original resources will be maintained on the companion website. Readers should confirm that they are consulting the most recent version of a resource before adopting or adapting particular language or practices.

Patterns across jurisdictions

The examples differ in authority, purpose, legal context, and level of detail, but several patterns recur:

- Guidance is often treated as living and subject to scheduled review.
- Several resources separate board or system commitments from administrative procedures, approved-tool lists, and classroom expectations.
- AI use is expected to serve a defined educational or operational purpose.
- Approval processes consider privacy, security, accessibility, age requirements, data practices, and potential consequences.
- Students and staff are expected to disclose relevant AI use and critically verify AI-generated content.
- Human beings retain responsibility for assessment and other decisions that may significantly affect students or staff.
- Professional learning is treated as a condition of responsible implementation.
- Equity, access, cultural context, Indigenous knowledge, and community communication are treated as governance concerns.
- Several Canadian resources caution against using AI-detection results as conclusive evidence of misconduct.
- Consultation, limited pilots, feedback, and monitoring are used to inform implementation and revision. These activities do not, by themselves, establish that a tool or practice is effective.

These patterns do not amount to a single accepted model of AI governance. They identify areas of emerging agreement while demonstrating the importance of local law, policy, curricula, relationships, technical capacity, and educational priorities. Saskatchewan boards can use these resources to ask better questions and compare approaches without assuming that any one framework can be adopted unchanged.

8. Keeping This Policy Advisory Current

Resolution AGM 25-04 calls for a review process so that this Advisory remains current as technologies evolve (SSBA, 2025a). The SSBA will update the Advisory as needed in light of developments in technology, law, evidence (including findings from Saskatchewan implementation), and sector experience. The timing, scope, participants, and form of each review remain at the SSBA's discretion.

A review may confirm that the current edition remains appropriate, identify targeted corrections or additions, or lead to a revised edition. Time-sensitive information and practical resources may be updated through the companion website between formal editions. Version numbers, publication dates, and brief change records can help boards identify the current guidance and understand substantive revisions.

Reviews should draw on findings from Saskatchewan divisions about learning, well-being, workload, equity, and unintended consequences. The SSBA could invite divisions to share these findings through sector consultation and use them to inform revisions and companion resources, without sharing identifiable student or employee information.

References

Citations in this document follow an author-date style. Each entry below includes sufficient detail to locate the source without a hyperlink; URLs are provided as a convenience and were current at the time of writing.

Alberta School Boards Association. (2024). Artificial intelligence policy guidance (Version 1.0). Edmonton, AB: ASBA. https://www.asba.ab.ca/uploads/files/Documents/Artificial%20Intelligence/ASBA%20AI%20Policy%20Guidance.pdf

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. Proceedings of the National Academy of Sciences, 122(26), Article e2422633122. https://doi.org/10.1073/pnas.2422633122

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, 1st Sess, 45th Parl, 2026 (first reading 10 June 2026). https://www.parl.ca/DocumentViewer/en/45-1/bill/C-34/first-reading

British Columbia Ministry of Education and Child Care. (n.d.). Considerations for using AI tools in K-12 schools. Victoria, BC: Government of British Columbia. https://www2.gov.bc.ca/assets/gov/education/administration/kindergarten-to-grade-12/ai-in-education/considerations-for-using-ai-tools-in-k-12-schools.pdf

Canada. *Copyright Act*, RSC 1985, c C-42.

Canadian Centre for Cyber Security. (2024). Recommended cyber security contract clauses for cloud services (ITSM.50.104). Communications Security Establishment Canada. https://www.cyber.gc.ca/en/guidance/recommended-cyber-security-contract-clauses-cloud-services-itsm50104

Canadian Centre for Cyber Security. (2026a). Ransomware playbook (ITSM.00.099). Communications Security Establishment Canada. https://www.cyber.gc.ca/en/guidance/ransomware-playbook-itsm00099

Canadian Centre for Cyber Security. (2026b). Social engineering (ITSAP.00.166). Communications Security Establishment Canada. https://www.cyber.gc.ca/en/guidance/social-engineering-itsap00166

Canadian Radio-television and Telecommunications Commission. (2025). Canadian Telecommunications Market Report 2025. Ottawa, ON: CRTC. https://crtc.gc.ca/eng/publications/reports/PolicyMonitoring/2025/ctmr.htm

Coglianese, C., & Crum, C. R. (2025). Leashes, not guardrails: A management-based approach to artificial intelligence risk regulation. Risk Analysis, 45(12), 4397–4407. https://doi.org/10.1111/risa.70020

Consortium for School Networking (CoSN), & Council of the Great City Schools (CGCS). (2024). CoSN/CGCS K-12 Gen AI maturity tool: Organizational assessment, readiness, and planning tool. https://www.cosn.org/the-gen-ai-maturity-model-organizational-assessment-readiness-and-planning-tool/

Council of Ministers of Education, Canada. (n.d.). Copyright: Overview and fair dealing resources. https://cmec.ca/736/Copyright.html

Couros, A. (2026). AI in education policy framework: WAG consultation survey: Findings and analysis [Unpublished internal report prepared for the Saskatchewan School Boards Association AI Working Advisory Group].

Couros, A., & Hildebrandt, K. (2015). Digital citizenship education in Saskatchewan schools: A policy planning guide for school divisions and schools to implement digital citizenship education from Kindergarten to Grade 12. Saskatchewan Ministry of Education. https://publications.saskatchewan.ca/api/v1/products/74447/formats/83322/download

De Simone, M. E., Tiberti, F. H., Barron Rodriguez, M., Manolio, F., Mosuro, W., & Dikoru, E. J. (2025). From chalkboards to chatbots: Evaluating the impact of generative AI on learning outcomes in Nigeria (Policy Research Working Paper No. 11125). World Bank. https://doi.org/10.1596/1813-9450-11125

EdCan Network. (2025). Artificial intelligence in K-12 education: Frameworks and guidance. Toronto, ON: EdCan Network. https://www.edcan.ca/wp-content/uploads/Ai-in-K-12-Frameworks.pdf

Educational Collaborative Network of Ontario (ECNO) & Ontario Association of School Business Officials (OASBO). (2025). Guidelines for responsible use of generative artificial intelligence (Version 1.3). https://www.ecno.org/wp-content/uploads/2025/04/English-Generative-AI-Guidelines-v1.3.pdf

Educational Collaborative Network of Ontario (ECNO) & Ontario Association of School Business Officials (OASBO). (2026). Artificial intelligence accountability framework. https://www.ecno.org/wp-content/uploads/2026/02/Artificial-Intelligence-Accountability-Framework-ENG-21226.pdf

European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

Federal, Provincial and Territorial Privacy Commissioners and Ombuds with Responsibility for Privacy Oversight. (2025). Protecting the privacy of children and youth through responsible use of educational technologies in the classroom. https://www.priv.gc.ca/en/about-the-opc/what-we-do/provincial-and-territorial-collaboration/joint-resolutions-with-provinces-and-territories/res_20251008_edtech/

Fesler, L., Martinez, J., Agnew, C., & Loeb, S. (2026). *The evidence base on AI in K–12: A 2026 review*. AI Hub for Education, SCALE Initiative, Stanford University. https://scale.stanford.edu/sites/default/files/The%20Evidence%20Base%20on%20AI%20in%20K-12%20Report.pdf

First Nations Information Governance Centre. (n.d.). The First Nations principles of OCAP®. Ottawa, ON: FNIGC. https://fnigc.ca/ocap-training/

Government of Canada. (2026). Canada's national artificial intelligence strategy: AI for all. Ottawa, ON: Innovation, Science and Economic Development Canada. https://ised-isde.canada.ca/site/ised/en/canadas-national-artificial-intelligence-strategy-ai-all

Government of Saskatchewan. (2024, August 6). New provincial measures to limit cell phones in Saskatchewan's K-12 classrooms [News release]. Regina, SK: Government of Saskatchewan. https://www.saskatchewan.ca/government/news-and-media/2024/august/06/new-provicial-measures-to-limit-cell-phones-in-saskatchewans-k-12-classrooms

Halton District School Board. (2025). Generative AI: Staff and student use guidelines. Burlington, ON: HDSB. https://www.myhdsb.ca/media/bhtc220m/generative-ai-staff-and-student-use-guidelines.pdf

Innovation, Science and Economic Development Canada. (2025). *What we heard report: Consultation on copyright in the age of generative artificial intelligence.* Government of Canada. https://ised-isde.canada.ca/site/publications/en/publications/what-we-heard-report-consultation-copyright-age-generative-artificial-intelligence

Innovation, Science and Economic Development Canada. (2026a, March 12). *Audio-visual recordings and privacy*. Government of Canada. https://ised-isde.canada.ca/site/ised/en/atip-services/references/audio-visual-recordings-and-privacy

Innovation, Science and Economic Development Canada. (2026b, June 15). *Government of Canada tables new legislation to protect children's data, strengthen privacy and build trust in the digital economy*. Government of Canada. https://www.canada.ca/en/innovation-science-economic-development/news/2026/06/government-of-canada-tables-new-legislation-to-protect-childrens-data-strengthen-privacy-and-build-trust-in-the-digital-economy.html

International Energy Agency (IEA). (2025). Energy and AI. Paris: IEA. https://www.iea.org/reports/energy-and-ai

Jiang, Y., Hao, J., Fauss, M., & Li, C. (2024). Detecting ChatGPT-generated essays in a large-scale writing assessment: Is there a bias against non-native English speakers? Computers & Education, 217, Article 105070. https://doi.org/10.1016/j.compedu.2024.105070

Langley School District (School District No. 35). (n.d.). SD35 AI guidelines: Use of artificial intelligence in Langley schools. Retrieved August 24, 2026, from https://www.sd35.bc.ca/instructionalservices/sd35-ai-guidelines

Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. Patterns, 4(7), Article 100779. https://doi.org/10.1016/j.patter.2023.100779

Moffatt v. Air Canada, 2024 BCCRT 149. https://www.canlii.org/en/bc/bccrt/doc/2024/2024bccrt149/2024bccrt149.html

National Cyber Security Centre. (2026). Managing the cyber risk of agentic AI. https://www.ncsc.gov.uk/blogs/managing-the-cyber-risk-of-agentic-ai

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). https://doi.org/10.6028/NIST.AI.100-1

Office of the Privacy Commissioner of Canada. (2024). Wearable devices and your privacy. https://www.priv.gc.ca/en/privacy-topics/technology/mobile-and-digital-devices/digital-devices/02_05_d_73_wd/

Office of the Privacy Commissioner of Canada, and international data protection authorities. (2026). Joint statement on AI-generated imagery and the protection of privacy. https://www.priv.gc.ca/en/opc-news/speeches-and-statements/2026/js-dc_20260223_ai/

Office of the Privacy Commissioner of Canada, and provincial and territorial privacy authorities. (2023). Principles for responsible, trustworthy and privacy-protective generative AI technologies. Ottawa, ON: OPC. https://www.priv.gc.ca/en/privacy-topics/technology/artificial-intelligence/gd_principles_ai/

Office of the Saskatchewan Information and Privacy Commissioner. (n.d.). Artificial intelligence (AI) privacy guidance [Resource collection, including materials on children's privacy and AI]. Regina, SK: OIPC. Retrieved October 20, 2026, from https://oipc.sk.ca/resources/ai-privacy-guidance/

Organisation for Economic Co-operation and Development (OECD). (2024). Explanatory memorandum on the updated OECD definition of an AI system. OECD Artificial Intelligence Papers, No. 8. OECD Publishing. https://doi.org/10.1787/623da898-en

Organisation for Economic Co-operation and Development (OECD). (2026). Reimagining teaching in an accelerating world. International Summit on the Teaching Profession, OECD Publishing. https://doi.org/10.1787/d0edfe8c-en

Parliament of Canada. (2026a). Bill C-34, 45th Parliament, 1st session [LEGISinfo bill record]. Retrieved October 12, 2026, from https://www.parl.ca/legisinfo/en/bill/45-1/c-34

Parliament of Canada. (2026b). Bill C-36, 45th Parliament, 1st session [LEGISinfo bill record]. Retrieved October 12, 2026, from https://www.parl.ca/legisinfo/en/bill/45-1/c-36

Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2024). The Artificial Intelligence Assessment Scale (AIAS): A framework for ethical integration of generative AI in educational assessment. Journal of University Teaching and Learning Practice, 21(6). https://doi.org/10.53761/q3azde36

Public Service Commission of Canada. (2026). Artificial intelligence in the hiring process. Government of Canada. https://www.canada.ca/en/public-service-commission/services/appointment-framework/guides-tools-appointment-framework/ai-hiring-process.html

Saskatchewan Human Rights Commission. (2023). Equitable education for students with reading disabilities in Saskatchewan's K to 12 schools: A systemic investigation report. https://saskatchewanhumanrights.ca/wp-content/uploads/2023/09/EQUITABLE-EDUCATION-for-Students-Reading-Disabilities-Report-2023.pdf

Saskatchewan Ministry of Education. (2022). Supporting student assessment in Saskatchewan. Government of Saskatchewan. https://publications.saskatchewan.ca/api/v1/products/121267/formats/140386/download

Saskatchewan School Boards Association (SSBA). (2025a). *2025 annual general meeting: Final results*. Regina, SK: SSBA. https://saskschoolboards.ca/wp-content/uploads/2025-AGM-Results.pdf

Saskatchewan School Boards Association (SSBA). (2025b). *2025 proposed bylaw amendments and resolutions (revised)*. Regina, SK: SSBA. https://saskschoolboards.ca/wp-content/uploads/2025-Proposed-Bylaw-Amendments-and-Resolutions-Revised.pdf

Saskatchewan School Boards Association (SSBA). (2026). About us. Regina, SK: SSBA. https://saskschoolboards.ca/about-us

Saskatchewan. The Education Act, 1995, SS 1995, c E-0.2. Regina, SK: Queen's Printer.

Saskatchewan. The Local Authority Freedom of Information and Protection of Privacy Act, SS 1990-91, c L-27.1. Regina, SK: Queen's Printer.

Saskatchewan. The Local Authority Freedom of Information and Protection of Privacy Regulations, RRS c L-27.1 Reg 1.

Saskatchewan. The Privacy Act, RSS 1978, c P-24.

Saskatchewan. The Saskatchewan Human Rights Code, 2018, SS 2018, c S-24.2. https://saskatchewanhumanrights.ca/wp-content/uploads/2020/03/S24-2.pdf

UK Government Digital Sustainability Alliance. (2025). Water use in AI and data centres. Government Digital Service. https://assets.publishing.service.gov.uk/media/688cb407dc6688ed50878367/Water_use_in_data_centre_and_AI_report.pdf

UN Committee on the Rights of the Child. (2021). General comment No. 25 (2021) on children's rights in relation to the digital environment (CRC/C/GC/25). Geneva: United Nations. https://docs.un.org/en/CRC/C/GC/25

UNESCO. (2021). Recommendation on the Ethics of Artificial Intelligence. Paris: UNESCO. https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence

UNESCO. (2023). Guidance for generative AI in education and research. Paris: UNESCO. https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research

UNESCO. (2024a). AI competency framework for students. Paris: UNESCO. https://www.unesco.org/en/articles/ai-competency-framework-students

UNESCO. (2024b). AI competency framework for teachers. Paris: UNESCO. https://www.unesco.org/en/articles/ai-competency-framework-teachers

UNICEF. (2025). Guidance on AI and children (Version 3.0). UNICEF Innocenti. https://www.unicef.org/innocenti/reports/policy-guidance-ai-children

United Nations. (1989). Convention on the Rights of the Child. New York: United Nations. https://www.ohchr.org/sites/default/files/crc.pdf

United Nations. (2007). United Nations Declaration on the Rights of Indigenous Peoples (A/RES/61/295). New York: United Nations. https://docs.un.org/en/A/RES/61/295

University of Saskatchewan. (2025). *AI literacy framework*. https://ai.usask.ca/ai-literacy-framework.php

Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P., & Waddington, L. (2023). Testing of detection tools for AI-generated text. International Journal for Educational Integrity, 19, Article 26. https://doi.org/10.1007/s40979-023-00146-z

APPENDIX A

Board Discussion Questions and Sample Policy Statements

This appendix contains the board discussion questions and core sample policy statements associated with Sections 6.1 through 6.7. It is organized as a set of pull-out sheets for board discussion and local policy development. Each subsection includes a question sheet, a copy-ready policy-statement sheet, and a drafting note explaining an important wording choice or policy distinction.

The questions and statements in this appendix are the authoritative copies used by the Advisory. Expanded or alternative statement options are not included and may be considered separately for the companion website after review.

About the statement codes: The short codes are reference identifiers used to locate, discuss, and select sample statements. The same codes are used in the companion website's Policy Generator so that statements can be matched reliably between the Advisory and the online tool. They are not proposed policy numbers and do not need to be retained when wording is incorporated into a board's policy framework.

Suggested use

- Print only the domains scheduled for discussion.
- Use the question sheet to identify policy choices, evidence, and follow-up work. In each domain, ask how the proposed use supports student learning or the conditions for learning, strengthens professional judgement, and sustains public trust. Ask administration to explain the expected benefit, the evidence supporting it, and how the division will check whether that benefit occurs.
- Copy selected policy statements into a local drafting document and adapt them to the division's existing policy architecture.
- Read selected statements together with the corresponding background in Section 6 before adoption.

The sample policy statements do not constitute legal advice. Boards should review proposed policies against applicable law and obtain legal advice where needed.

APPENDIX A | BOARD DISCUSSION QUESTIONS

A.1 AI literacy, professional learning, and well-being

Common test: Support student learning or the conditions for learning, strengthen professional judgement, and sustain public trust.

Use this sheet to identify what the board can verify, what requires clarification, and what should be addressed through policy or direction to administration.

	A.1 | Questions for board discussion
1	Does the division have a coordinated plan for role-appropriate AI literacy learning and support for students, staff, division leaders, and board members, consistent with provincial curricula and professional responsibilities?
2	Do those expectations address how AI works and where it can fail, as well as verification, privacy, responsible use, and whether AI supports rather than substitutes for learning or work?
3	Is professional learning sustained, role-specific, current, adequately resourced, and inclusive of the division's responsibilities concerning Indigenous authority and knowledge governance?
4	Are expectations for student access appropriate to age, developmental readiness, educational purpose, supervision, privacy, and evidence of benefit?
5	How will the division protect student and staff well-being, human relationships, and professional judgement as AI use expands?
6	How will the division provide equitable access to AI literacy learning and support across schools and communities with different levels of connectivity, staffing, and local capacity?

NOTES, EVIDENCE, OR FOLLOW-UP

APPENDIX A | COPY-READY POLICY STATEMENTS

A.1 AI literacy, professional learning, and well-being

Core sample policy statements associated with Section 6. Select and adapt statements within the context of the division's existing policy framework.

Code	A.1 | Core sample policy statements
LIT-01	The Board recognizes AI literacy as a core capacity for students, staff, division leaders, and board members.
LIT-02	The Board expects division administration to support role-appropriate AI literacy concerning how AI systems produce outputs and where they may fail, verification, privacy, responsible use, approved systems, and human judgement, consistent with provincial curricula and professional responsibilities.
LIT-03	The Board expects AI use to protect and strengthen critical thinking, foundational knowledge and skills, independent work, human relationships, and professional judgement.
LIT-04	The Board expects division administration to provide sustained, role-appropriate professional learning for teachers, administrators, support staff, and division leaders. The Board will also support continuing AI literacy development for board members.
LIT-05	The Board expects student access to AI tools to reflect age, developmental readiness, educational purpose, supervision, privacy, and evidence of benefit. Specific requirements shall be established in administrative procedure.
LIT-06	The Board expects division administration to involve staff in AI implementation and to address the effects of AI on workload, professional judgement, relationships, and well-being.
LIT-07	The Board expects division administration to provide equitable access to AI literacy learning and support across schools and communities, with attention to differences in connectivity, staffing, and local capacity.

Drafting note

- **LIT-05.** This statement identifies the factors the Board expects division administration to address while leaving specific age limits, supervision, consent, and system-approval requirements to administrative procedure.
- **Relationship to Section 6.5.** Section 6.1 addresses the literacy, developmental, relational, and well-being considerations associated with conversational AI. Section 6.5 establishes the approval boundaries, monitoring expectations, and safety safeguards for its use with students.

APPENDIX A | BOARD DISCUSSION QUESTIONS

A.2 Academic integrity and assessment

Common test: Support student learning or the conditions for learning, strengthen professional judgement, and sustain public trust.

Use this sheet to identify what the board can verify, what requires clarification, and what should be addressed through policy or direction to administration.

	A.2 | Questions for board discussion
1	Has the board established division-wide principles for communicating permitted AI use and disclosure requirements, addressing academic-integrity concerns fairly, using approved systems, and preserving accountable human judgement?
2	Does administration provide common language and guidance that distinguishes division-wide expectations from school-, course-, or task-specific requirements and preserves teachers' professional judgement?
3	Do academic-integrity procedures protect confidentiality, require relevant evidence, and provide students with a meaningful opportunity to respond?
4	Has the board determined that the division will not procure AI-text detection tools for the purpose of determining whether a student engaged in academic misconduct?
5	Does administration provide the procedures, professional learning, and resources needed to implement these expectations consistently and fairly?
6	Is it clear that assessment and evaluation remain the professional responsibility of teachers and that no summative evaluation or final decision about student achievement may be made solely by an automated system?
7	Are students informed when AI materially contributes to individualized feedback and whether that feedback was reviewed by a teacher or delivered directly by an approved instructional system?

NOTES, EVIDENCE, OR FOLLOW-UP

APPENDIX A | COPY-READY POLICY STATEMENTS

A.2 Academic integrity and assessment

Core sample policy statements associated with Section 6. Select and adapt statements within the context of the division's existing policy framework.

Code	A.2 | Core sample policy statements
INT-01	The Board expects division administration to establish common language and guidance for communicating division-wide expectations for AI use and to ensure that applicable school-, course-, or task-specific requirements are made clear to students, while preserving teachers' professional judgement.
INT-02	Students shall follow the requirements communicated for AI use and make any disclosure required for the task. AI use contrary to those requirements, including submitting AI-generated content as the student's own when that use is not permitted or disclosed as required, shall be addressed through established academic-integrity procedures.
INT-03	The division shall not procure an AI-text detector for the purpose of determining whether a student engaged in academic misconduct. A detector score available through an otherwise approved system shall not be treated as a finding of misconduct or used as the sole basis for an allegation or decision.
INT-04	Academic-integrity decisions shall follow established division procedures, be based on relevant evidence, protect confidentiality, and provide the student with a meaningful opportunity to respond.
INT-05	Assessment and evaluation of student learning remain the professional responsibility of teachers. Approved AI systems may assist under conditions established by administration, but no summative evaluation or final decision about student achievement shall be made solely by an automated system.
INT-06	When AI materially contributes to individualized feedback, the student shall be informed that AI assistance was used and whether the feedback was reviewed by a teacher or delivered directly by an approved instructional system.
INT-07	AI systems used in assessment, feedback, evaluation, or academic-integrity processes shall be approved by division administration for the intended purpose and shall comply with the division's privacy, security, information-handling, and intellectual-property requirements.

Drafting note

- **INT-01.** This statement establishes a communication chain. Division administration provides common division-wide language, schools make general expectations accessible, and teachers explain how those expectations apply to particular learning and assessment tasks.
- **INT-03.** This statement does not prevent staff from addressing a reasonable concern about unauthorized AI use. It recognizes that the uncertain and context-dependent performance of detector tools does not justify procuring them to determine misconduct. A detector score available through another approved system cannot replace relevant evidence and a fair process.
- **INT-05 and INT-06.** These statements permit carefully governed AI assistance while keeping responsibility with the teacher. The teacher reviews the AI contribution, may correct or reject it, makes the final professional judgement, and remains accountable for the result. Proportionate disclosure is required when AI materially contributes to individualized feedback, including whether the feedback was reviewed by a teacher or delivered directly by an approved instructional system. Administrative procedures may establish the conditions under which approved systems provide formative assistance.
- **Division of responsibility.** The Board establishes division-wide principles and accountability. Division administration implements them through procedures, system approval, professional learning, and shared resources. Teachers retain professional responsibility for learning and assessment design, the evidence appropriate to particular tasks, and evaluation of student achievement within provincially authorized curricula, board policy, and division guidance.
- **Related sections.** The general professional-learning obligation is addressed in Section 6.1. Privacy and system-approval requirements are addressed in Sections 6.3 and 6.4. Section 6.5 establishes the broader rules for prohibited and restricted uses, including fully automated consequential decisions and meaningful human responsibility.

APPENDIX A | BOARD DISCUSSION QUESTIONS

A.3 Privacy, security, and data governance

Common test: Support student learning or the conditions for learning, strengthen professional judgement, and sustain public trust.

Use this sheet to identify what the board can verify, what requires clarification, and what should be addressed through policy or direction to administration.

	A.3 | Questions for board discussion
1	Does board policy make clear that LA FOIP applies whenever the division uses AI systems or AI-enabled features to handle personal information, and does it assign responsibility for implementation and oversight?
2	Does the division's approved-use record identify which AI systems handle personal information, their approved purposes, responsible owners, providers, the categories of information involved, and where that information is stored or processed?
3	Does administration complete and document a privacy review before approving a new or materially changed use of an AI system that handles personal information, with more extensive review when the information is sensitive or the potential consequences are significant?
4	Where a provider handles personal information through an AI system or AI-enabled feature, does the written agreement meet LA FOIP and regulatory requirements, expressly cover the AI-enabled functions, and specify whether personal information may be used for model training or service improvement and the purposes and limits of any permitted use?
5	Are approved AI systems incorporated into existing access-control, cybersecurity, incident-response, business-continuity, records-management, and enterprise-risk practices, and do established reporting processes provide the board with information about significant risks, incidents, and corrective actions?
6	Where information has collective First Nations or Métis dimensions or includes Indigenous knowledge or community-contributed material, has the division worked with the relevant rights-holding First Nation, Métis government or community, or First Nations education authority to identify the applicable governance requirements?

NOTES, EVIDENCE, OR FOLLOW-UP

APPENDIX A | COPY-READY POLICY STATEMENTS

A.3 Privacy, security, and data governance

Core sample policy statements associated with Section 6. Select and adapt statements within the context of the division's existing policy framework.

Code	A.3 | Core sample policy statements
PRIV-01	*The Local Authority Freedom of Information and Protection of Privacy Act* applies whenever the division's use of an AI system or AI-enabled feature involves personal information. The division shall collect, use, disclose, retain, and protect that information in accordance with the Act and division policy.
PRIV-02	Personal information shall not be entered into an AI system unless division administration has assessed and approved the system and use for that purpose. The approved use shall be limited to the information reasonably necessary for its stated purpose and shall use division-managed accounts, settings, and workflows where required.
PRIV-03	Before administration approves a new or materially changed use of an AI system that handles personal information, it shall complete and document a privacy review using the division's established process. The depth of review shall reflect the sensitivity and scale of the information, the people affected, and the potential consequences.
PRIV-04	Before personal information is disclosed to an information management service provider, the division shall enter into a written agreement that meets section 23.2 of LA FOIP and section 8.2 of its Regulations. Where the service includes an AI system or AI-enabled feature, the agreement shall expressly cover those functions and specify whether personal information may be used for model training or service improvement and the purposes and limits of any permitted use.
PRIV-05	Division administration shall conduct a renewed privacy and security review of an approved AI use when its purpose, information practices, model, integrations, subcontractors, or other relevant conditions materially change.
PRIV-06	The division shall incorporate AI systems into its established access-control, cybersecurity, incident-response, business-continuity, records-management, and enterprise-risk practices. These practices shall address AI-enabled threats, and significant risks, incidents, and corrective actions shall be reported to the Board in accordance with established reporting requirements.
PRIV-07	Before Indigenous knowledge or information with collective First Nations or Métis dimensions is used with AI, the division shall work with the relevant rights-holding First Nation, Métis government or community, or First Nations education authority to identify and respect the applicable authority, permissions, and data-governance requirements. OCAP® shall be applied as determined by the relevant First Nation and shall not be treated as a generic framework for all Indigenous Peoples.

Drafting note

- **PRIV-02.** Approval is specific to both the AI system and the proposed use. Approval for work involving non-confidential information does not authorize the same system to process personal information or connect to division records.
- **PRIV-03.** This statement uses the division's established privacy-review process rather than requiring a separate AI procedure. The depth of review should reflect the sensitivity and scale of the information, the people affected, and the potential consequences. A formal privacy impact assessment may be appropriate for a higher-risk use, but it is not required for every AI use.
- **PRIV-04.** Section 23.2 of LA FOIP and section 8.2 of its Regulations establish the legal requirements for provider agreements. The agreement should apply those requirements expressly to AI-enabled functions, rather than duplicate them in a separate list. Provisions governing model training or service improvement should specify permitted uses and limits. These AI-specific details are recommended contractual safeguards. Section 6.4 addresses the broader procurement process, provider relationship, renewal, intellectual property, and exit planning.
- **PRIV-06.** This statement uses existing division governance structures rather than requiring a separate AI cybersecurity or risk system. Division administration remains responsible for technical controls and incident management. Board reporting should be proportionate to the significance of the risk or incident.
- **PRIV-07.** This statement applies when information has collective First Nations or Métis dimensions or includes Indigenous knowledge or community-contributed material. It does not imply that every ordinary student record engages collective data governance. LA FOIP continues to govern personal information, while applicable Indigenous authority and governance requirements may impose additional responsibilities.
- **Relationship to Section 6.5.** Section 6.3 addresses whether information may be collected, entered, processed, disclosed, retained, secured, and governed. Section 6.5 addresses prohibited and restricted uses and their procedural safeguards. A use may satisfy the privacy requirements in Section 6.3 and still be prohibited or restricted for another reason.

APPENDIX A | BOARD DISCUSSION QUESTIONS

A.4 AI tools, procurement, and intellectual property

Common test: Support student learning or the conditions for learning, strengthen professional judgement, and sustain public trust.

Use this sheet to identify what the board can verify, what requires clarification, and what should be addressed through policy or direction to administration.

	A.4 | Questions for board discussion
1	Does the division have one accountable pathway for acquiring, piloting, implementing, renewing, and retiring AI systems, including free tools, trials, and AI features added to software the division already licenses?
2	When administration reviews a proposed AI system or use through the division's technology and procurement process, does it document the purpose and expected benefit, the division's ability to implement and sustain the system, and the applicable legal, policy, and contractual requirements?
3	Do provider agreements address service and support, accessibility, who may use prompts, uploaded files, outputs, and other material, oversight, material changes, pricing, portability, transition, termination, and long-term dependency?
4	Do staff have clear direction concerning protected material submitted to AI systems and the review, attribution, permissions, and disclosure required for AI-assisted official materials?
5	Does the division protect Indigenous authority when proposed systems or agreements may involve Indigenous knowledge, cultural expressions, or community-contributed material?
6	Before requiring direct student use of a particular AI system, has the division addressed age and account requirements, accessibility, accommodation, privacy, consent, supervision, support, and reasonable alternatives?

NOTES, EVIDENCE, OR FOLLOW-UP

APPENDIX A | COPY-READY POLICY STATEMENTS

A.4 AI tools, procurement, and intellectual property

Core sample policy statements associated with Section 6. Select and adapt statements within the context of the division's existing policy framework.

Code	A.4 | Core sample policy statements
TOOL-01	The Board expects division administration to maintain a documented process for acquiring, piloting, implementing, renewing, and retiring AI systems and AI-enabled features, including free tools, trials, and features added to software the division already licenses.
TOOL-02	Before approving an AI system or proposed use, administration shall document its purpose and expected benefit, implementation requirements and total cost, and the applicable legal, policy, and contractual requirements.
TOOL-03	When administration approves an AI system, it shall record the system and service version, permitted purposes, accountable administrator or unit, authorized users or groups, categories of information that may be handled, conditions of use, and required reassessment. Approval for one purpose does not authorize other uses of the same system.
TOOL-04	A pilot shall comply with the privacy, security, accessibility, restricted-use, contractual, and other division requirements applicable to the proposed use. Administration shall define the pilot's purpose, duration, safeguards, accountable administrator or unit, and evaluation criteria.
TOOL-05	Agreements for AI systems shall address service and support, accessibility, who owns or may use prompts, uploaded files, outputs, and other material, oversight, material changes, pricing, portability, transition, termination, and continuing obligations.
TOOL-06	Before renewing an agreement for an AI system, division administration shall review continuing need, performance, cost, dependency, material changes, and available alternatives.
TOOL-07	Staff shall not accept provider terms for an AI system outside the division's authorized procurement and contracting process.
TOOL-08	Protected or confidential material, including student-created and staff-created work, shall be submitted to AI systems only through practices authorized by administration and consistent with applicable law, ownership, licences, agreements, permissions, and division policy. Division procedures shall identify who may authorize submission or reuse of the material.
TOOL-09	Staff and administrators shall review AI-assisted official materials before publication or distribution and remain responsible for accuracy, attribution, permissions, confidentiality, and compliance. Division procedures shall establish when material AI assistance is to be disclosed.

Code	A.4 | Core sample policy statements
TOOL-10	Division approval shall not replace any direction or authorization required from the relevant First Nation, Métis government or community, or First Nations education authority before Indigenous knowledge, language, stories, cultural expressions, or community-contributed material is submitted to an AI system or used for training, service improvement, product development, or another purpose.
TOOL-11	Direct student use of a particular AI system shall be required only where the system and use have been approved for the relevant students and purpose and applicable age, accessibility, accommodation, privacy, consent, supervision, and support requirements have been addressed. The division shall provide a reasonable alternative where required.

Drafting note

- **TOOL-01 through TOOL-04.** These statements establish one pathway for acquiring, piloting, implementing, renewing, and retiring AI systems. The Board establishes the requirement. Division administration determines the workflow and responsibilities. The same record should support procurement and any additional review required under other policies.
- **TOOL-05 through TOOL-07.** These statements address the broader relationship with an AI provider. The statutory agreement and information-handling requirements that apply when a provider handles personal information are addressed in Section 6.3.
- **TOOL-08 and TOOL-09.** These statements distinguish material submitted to an AI system from material produced with AI assistance. Disclosure of AI assistance is not necessarily the same as citation or attribution. Requirements applying to student learning and assessment tasks are addressed in Section 6.2.
- **TOOL-10.** Division approval cannot replace any applicable Indigenous authority. The broader principles governing Indigenous knowledge and data are addressed in Sections 5 and 6.3.
- **TOOL-11.** This statement applies when students are required to use a particular AI system directly. It does not establish a universal right to opt out of every AI-enabled feature.
- **Related sections.** Section 6.5 determines whether a proposed use is prohibited or restricted and establishes safeguards for higher-risk uses. Section 6.6 addresses division-wide patterns in the distribution of access, benefit, cost, workload, risk, and burden. Section 6.7 applies these requirements to procurement and other administrative and business functions.

APPENDIX A | BOARD DISCUSSION QUESTIONS

A.5 Prohibited and restricted uses

Common test: Support student learning or the conditions for learning, strengthen professional judgement, and sustain public trust.

Use this sheet to identify what the board can verify, what requires clarification, and what should be addressed through policy or direction to administration.

	A.5 | Questions for board discussion
1	Are prohibited and restricted uses clearly distinguished, stated in plain language, and defined by conduct and purpose rather than particular products?
2	Does policy respect the authority of First Nations and Métis rights holders where AI use involves Indigenous language, ceremony, ceremonial protocol, sacred knowledge, or community-held information?
3	Does policy prohibit school-related intimate or sexualized synthetic depictions and harmful impersonation, with support and response processes for those affected?
4	Do staff and students have clear boundaries for uploading protected material to AI systems, and are copyright, licensing, ownership, attribution, and publication review addressed elsewhere in policy and administrative procedure?
5	Do consequential decisions remain under meaningful human responsibility, and are AI-assisted learning analytics subject to reliability, bias, explanation, correction, and review requirements?
6	Are biometric identification, e-proctoring, behavioural inference, surveillance, and AI-enabled wearables subject to clear and proportionate limits that preserve approved accommodations?
7	Are instructional chatbots, service chatbots, and companion-style systems distinguished, with appropriate requirements for privacy, accuracy, human contact, monitoring, and safety escalation?
8	Are AI agents limited to approved purposes, permissions, actions, and environments, with human authorization and a reliable means of stopping them?
9	Do responses to misuse remain educative and proportionate while ensuring that conduct causing harm or compromising protected material, systems, or information receives a stronger response?

APPENDIX A | COPY-READY POLICY STATEMENTS

A.5 Prohibited and restricted uses

Core sample policy statements associated with Section 6. Select and adapt statements within the context of the division's existing policy framework.

Code	A.5 | Core sample policy statements
USE-01	The division will not use AI to teach an Indigenous language or to generate, reproduce, or prepare material involving ceremony, ceremonial protocol, or sacred knowledge unless the relevant First Nation, Métis government or community, or First Nations education authority has directed or authorized the use. The division will engage relevant rights holders, Elders, and Knowledge Keepers in developing its practice in this area. Engagement and any required direction or authorization shall follow the protocols of the relevant nation or community for the particular knowledge and proposed use.
USE-02	The creation, solicitation, or distribution of AI-generated or AI-altered intimate or sexualized images, audio, or video is prohibited when the conduct involves a student, uses division systems, occurs through a school-related activity, involves a division employee acting in a professional role, or affects safety, employment, or the learning environment. The conduct shall be addressed with the seriousness of other forms of abuse, with support for those harmed.
USE-03	AI shall not be used to impersonate an identifiable person in a manner intended or reasonably likely to deceive, harass, exploit, defraud, or cause harm.
USE-04	Copyrighted, licensed, student-created, staff-created, and community-contributed material shall be submitted to AI systems only through approved practices consistent with applicable law, licences, agreements, and division policy. AI shall not be used knowingly to evade access or licence restrictions or remove or falsify attribution.
USE-05	No automated system shall make a final decision that significantly affects a student, staff member, or other person. A responsible person with the information, time, and authority to reach a different conclusion shall remain accountable for the decision.
USE-06	AI-assisted analytics concerning identifiable students shall be limited to approved purposes and shall not be the sole basis for a consequential decision. Administrative procedures shall address reliability, bias, explanation, correction, human review, and accountability.
USE-07	Division administration shall specifically approve any use of biometric identification or verification or e-proctoring involving biometric or behavioural monitoring based on necessity, proportionality, accuracy, security, privacy, accessibility, and reasonable alternatives. Automated inferences about emotion, attention, honesty, or intent shall not be used as a basis for assessment, discipline, employment decisions, or routine monitoring.
USE-08	AI shall not be used for covert or continuous surveillance beyond what existing law and division policy expressly authorize for a defined purpose.

Code	A.5 | Core sample policy statements
USE-09	AI-enabled wearable devices shall not be used to record, stream, transcribe, identify, analyze, or transmit information about another person except as authorized by law and division policy. The division shall meet medical and accessibility accommodation needs through specifically approved functions or an effective alternative that protects the privacy and rights of others.
USE-10	AI-enabled wearable devices shall not be used to obtain assistance where that assistance is prohibited. The division shall meet medical and accessibility accommodation needs through specifically approved functions or an effective alternative that protects the privacy and rights of others.
USE-11	Student-facing instructional chatbots shall be limited to approved educational uses under age-appropriate conditions and shall clearly identify themselves as AI.
USE-12	Division administration shall specifically approve any service chatbot that accesses records, accepts identifiable information, recommends services, or handles sensitive matters. Administrative procedures shall address accuracy, privacy, security, record handling, logging, human contact, correction, and accountable ownership.
USE-13	Companion-style AI shall not be provided or represented as a substitute for counselling, crisis response, peer relationships, or trusted adult support. Where a conversational system uses automated safety alerts, administrative procedures shall define alert conditions, human responsibilities, response expectations, record access, and communication with students and families.
USE-14	AI agents capable of acting within division systems shall be limited to approved purposes, permissions, actions, and environments. Consequential actions require human authorization, activity logging, proportionate monitoring, and a reliable means of suspending the agent. AI shall not be used to bypass security controls or gain unauthorized access to systems, accounts, or information.
USE-15	Responses to student misuse of AI shall be educative and proportionate, except where conduct harms another person, threatens safety, deliberately compromises protected material, or compromises division systems or information. Such conduct shall be addressed through established processes, with support for those harmed.

Drafting note

- **USE-02.** This statement defines the connection between the conduct and the division so that board policy does not purport to govern private adult conduct unrelated to the division. Whether particular conduct engages criminal or civil law must be determined under the applicable legislation.
- **USE-04.** This statement addresses the AI-specific risk created when protected material is uploaded or disclosed to a provider. It does not relocate fair dealing, licensing, authorship, ownership, attribution, student and staff rights, provider terms, or publication review from Section 6.4.
- **USE-09 and USE-10.** Approval of an accommodation does not authorize every device or function. Restrictions must be based on an individualized assessment and must preserve effective accommodation, including through an alternative where necessary.
- **USE-11 through USE-13.** These statements distinguish instructional chatbots, service chatbots, and companion-style systems so that boards may adopt or adapt requirements for each separately. Section 6.1 addresses the literacy, developmental, relational, and well-being considerations associated with conversational systems. Section 6.5 establishes approval boundaries and safeguards. A system that combines conversational, analytical, or agentic functions should be governed according to its highest-risk function and the consequences of failure.

APPENDIX A | BOARD DISCUSSION QUESTIONS

A.6 Equity, accountability, and community engagement

Common test: Support student learning or the conditions for learning, strengthen professional judgement, and sustain public trust.

Use this sheet to identify what the board can verify, what requires clarification, and what should be addressed through policy or direction to administration.

	A.6 | Questions for board discussion
1	Can administration identify material differences among schools or affected groups in access to AI systems and service levels, benefits, costs, additional workloads, risks, or burdens, and are findings requiring action reported to the board?
2	Which differences in connectivity, equipment, staffing, accessibility, language, and local capacity require division support, an alternative approach, interdivisional collaboration, or provincial action?
3	How is responsibility for coordinating AI governance assigned, and do the people responsible have the expertise, organizational authority, information, and access to leadership required to act?
4	Does reporting enable the board to identify significant uses and changes, serious incidents and recurring concerns, costs and capacity pressures, evidence of benefits and limitations, material equity effects, and matters requiring board direction?
5	Can students, staff, families, and community members find accessible information about approved student-facing AI systems, their purposes, the general categories of information they handle, and routes for questions, correction, concerns, or review?
6	Are people informed when AI materially contributes to content, recommendations, or decisions that affect them and disclosure is relevant to trust, accountability, or their ability to question the result?
7	Are affected groups, including students, engaged early enough to influence consequential decisions, and does the division communicate what it heard and how the input affected the decision?
8	What requirements govern the board's own use of AI, including approved systems, confidentiality, review, disclosure, board member learning, and periodic self-evaluation?

NOTES, EVIDENCE, OR FOLLOW-UP

APPENDIX A | COPY-READY POLICY STATEMENTS

A.6 Equity, accountability, and community engagement

Core sample policy statements associated with Section 6. Select and adapt statements within the context of the division's existing policy framework.

Code	A.6 | Core sample policy statements
EQ-01	The Board expects division administration to monitor how access to AI systems and service levels, benefits, costs, workloads, risks, and burdens are distributed across schools, students, staff, and communities and to identify material inequities requiring action.
EQ-02	Division administration shall identify whether differences in connectivity, equipment, staffing, accessibility, language, or local capacity affect division-wide implementation and shall recommend necessary supports, alternatives, collaborative approaches, or requests for provincial action.
EQ-03	The Director of Education shall ensure that responsibility for coordinating AI governance is assigned at an organizational level with sufficient authority, access to relevant expertise and information, and a clear route for escalating significant matters.
EQ-04	The Board shall receive information on AI implementation at least annually, through existing reporting structures where appropriate. The information shall address significant uses and changes, serious incidents and recurring concerns, costs and capacity pressures, evidence of benefits and limitations, material equity and accessibility effects, corrective actions, and matters requiring Board direction. Matters requiring prompt attention shall be reported through established escalation processes without waiting for the annual reporting cycle.
EQ-05	The division shall maintain accessible public information about its AI policies, approved student-facing systems and purposes, the general categories of information those systems handle, user expectations, and routes for questions, correction, concerns, or review.
EQ-06	Division administration shall establish proportionate transparency requirements for staff and administrative uses of AI. Affected people shall be informed when AI materially contributes to content, recommendations, or decisions and knowledge of that assistance is relevant to trust, accountability, or the ability to question the result.
EQ-07	The division shall provide meaningful opportunities for affected students, staff, families, school community councils, and communities to inform consequential AI decisions. It shall communicate what was heard and how the input affected the decision.
EQ-08	Community engagement under this policy shall not replace any direction or authorization required from a First Nation, Métis government or community, or First Nations education authority.
EQ-09	The Board shall apply the division's AI governance expectations to its own work, including approved systems, confidentiality, review of AI-assisted content, human responsibility, disclosure of material AI assistance, board member learning, and periodic self-evaluation.

Drafting note

- **EQ-01 and EQ-02.** These statements concern division-wide patterns and cumulative effects. Section 6.1 addresses equitable opportunities to develop AI literacy. Section 6.4 addresses accessibility, implementation, and cost when a particular system or use is reviewed.
- **EQ-04.** This statement establishes an annual minimum for governance reporting. It does not delay serious incidents or material changes that require prompt Board attention. Division administration should use established reporting and escalation processes rather than create a separate channel for every AI matter.
- **EQ-05 and EQ-06.** These statements distinguish public information about division systems from disclosure that AI materially contributed to particular content, recommendations, or decisions. Detailed privacy notices and individual procedural safeguards are addressed through the applicable privacy and restricted-use processes.
- **EQ-07 and EQ-08.** These statements distinguish general community engagement from the direction or authorization that may be required from First Nations or Métis rights holders.
- **EQ-09.** This statement applies the division's AI governance expectations to the Board's own work. Section 6.7 applies the six policy domains to division administration and business functions.

APPENDIX A | BOARD DISCUSSION QUESTIONS

A.7 Applying the domains: division administration and business functions

Common test: Support student learning or the conditions for learning, strengthen professional judgement, and sustain public trust.

Use this sheet to identify what the board can verify, what requires clarification, and what should be addressed through policy or direction to administration.

	A.7 | Questions for board discussion
1	Which administrative and business functions use AI systems or AI-enabled features, and are those uses recorded with their purposes, accountable administrators or units, permitted users, information involved, conditions, and required reassessment?
2	Where AI informs a consequential administrative decision, can the responsible person understand and question the output, consider relevant evidence and circumstances, reach a different conclusion, and remain accountable for the decision?
3	Are recruitment and employment uses fair, accessible, explainable, and subject to clear expectations for candidate AI use, appropriate notice, meaningful human judgement, and routes for review?
4	Before approving a workplace use that may materially affect employee duties, workload, professional judgement, monitoring, or employment, does administration assess its purpose, expected benefit, reliability, possible bias or barriers, privacy, accessibility, proportionality, and less intrusive alternatives?
5	Before such a workplace use is implemented, how are its effects on employees assessed, how are affected staff and bargaining agents involved, and what professional learning, role clarification, or other support is provided?
6	Do financial and purchasing processes independently verify material information and sensitive requests, preserve delegated authority and segregation of duties, and protect confidential procurement material?
7	What requirements govern AI recording, transcription, summarization, retention, participant objections, official records, access-to-information searches, and use in confidential or in-camera meetings?
8	Are official communications, translations, and public-facing AI services reviewed for accuracy, accessibility, confidentiality, and safety, with proportionate disclosure and an accessible route to a person?
9	Do AI-assisted operational analyses identify their sources, assumptions, limitations, alternatives, and accountable decision-makers, and are physical systems protected by appropriate safety and human-intervention controls?

	A.7 | Questions for board discussion
10	Where AI agents can access information, communicate, initiate transactions, or act within division systems, are their permissions, human authorizations, activity records, monitoring, and means of suspension clearly defined?
11	Can administration identify the functions that must continue during a technology outage, the systems and providers on which they depend, how long each function can be unavailable before serious harm or disruption occurs, the order in which services must be restored, and the tested manual or alternate processes available during the disruption?
12	Does board reporting show whether administrative AI uses that may materially affect people, public funds, services, or essential operations are meeting their approved purposes after accounting for errors, verification, complaints, staff workload, full costs, provider dependence, and operational risk?

APPENDIX A | COPY-READY POLICY STATEMENTS

A.7 Applying the domains: division administration and business functions

Core sample policy statements associated with Section 6. Select and adapt statements within the context of the division's existing policy framework.

Code	A.7 | Core sample policy statements
OPS-01	The Board expects the division's requirements for approved use, privacy, security, purchasing, restricted uses, transparency, and accountability to apply wherever AI systems or AI-enabled features are used, including in administrative and business functions.
OPS-02	AI may assist with administrative work only where the system and use are approved for the intended purpose, information is handled through authorized practices, and a responsible person reviews and remains accountable for the resulting work. Approval for routine assistance does not authorize a consequential use.
OPS-03	Recruitment and employment practices involving AI shall remain fair, accessible, explainable, and subject to meaningful human judgement. The division shall establish clear expectations for candidate use of AI, and applicants and employees shall receive appropriate notice when AI materially influences a process affecting them.
OPS-04	A proposed workplace use of AI that may materially affect employee duties, workload, professional judgement, monitoring, or employment shall be approved only where its purpose and expected benefit are defined and administration has considered the system's reliability and limits, possible bias or barriers, privacy, accessibility, effects on employees, proportionality, and less intrusive alternatives.
OPS-05	Before an approved workplace use that may materially affect employee duties, workload, professional judgement, monitoring, or employment is implemented, division administration shall involve affected staff and bargaining agents as appropriate and identify required professional learning, role clarification, and other supports.
OPS-06	Financial and purchasing processes shall preserve delegated authority, segregation of duties, documentation, independent review, and auditability. Material or unusual requests, including changes to supplier or employee banking information, shall be verified through an established channel independent of the request. An email, text message, collaboration-platform message, displayed telephone number, voice recording, video call, or AI-generated document shall not by itself authorize a payment or change to financial information.

Code	A.7 | Core sample policy statements
OPS-07	AI recording, transcription, summarization, translation, or analysis shall be used in meetings only for a purpose approved by division administration and under conditions established in administrative procedure. Participants shall receive advance notice. Such functions shall not be used during an in-camera board session, whether online, hybrid, or in person, unless the Board has authorized the use and administration has confirmed that privacy, security, confidentiality, and records-management requirements have been addressed. Division administration may authorize the use during another protected meeting only after those requirements have been addressed. Approved minutes shall remain the official record of board and committee decisions.
OPS-08	AI-assisted official communications, reports, recommendations, and records shall be reviewed by an accountable staff member. AI assistance shall be disclosed when it materially contributes to the content and knowledge of that assistance is relevant to trust, accountability, or the ability of an affected person to question the result.
OPS-09	Public-facing AI systems shall be clearly identified, regularly checked for accuracy, and supported by an accessible route to a person. Where a system may receive urgent safety disclosures, its monitoring, escalation, response expectations, and emergency limitations shall be established and communicated.
OPS-10	AI-assisted analyses that may significantly affect schools or communities shall identify their sources, assumptions, limitations, and relevant alternatives. AI connected to physical infrastructure shall not independently alter a safety-critical system without defined authority, testing, monitoring, and a reliable means of human intervention.
OPS-11	AI agents shall operate only within an approved purpose, defined access and actions, documented human-authorization requirements, activity recording, proportionate monitoring, and a reliable means of suspension. They shall not independently authorize payments, make consequential decisions about people, change access permissions, delete official records, or communicate externally on behalf of the division.
OPS-12	The division shall identify the essential functions supported by AI-enabled or third-party systems and shall maintain and test incident-response, recovery, communication, and continuity arrangements, including service-restoration priorities and workable manual or alternate processes.
OPS-13	Administration shall evaluate administrative AI uses that may materially affect people, public funds, services, or essential operations for continuing need, performance, errors, verification requirements, complaints, staff workload, full cost, provider dependence, and operational risk and shall report material findings through established governance processes.

Drafting note

- **Purpose of Section 6.7.** Section 6.7 applies the six policy domains to administrative and business functions. It does not replace the professional-learning requirements in Section 6.1, the privacy and security requirements in Section 6.3, the approval and procurement requirements in Section 6.4, the prohibited and restricted-use requirements in Section 6.5, or the accountability, transparency, engagement, and reporting framework in Section 6.6.
- **OPS-02 through OPS-05.** OPS-02 makes approval specific to the proposed administrative purpose. OPS-03 addresses recruitment and employment decisions. OPS-04 and OPS-05 distinguish the criteria for approving a workplace use from the engagement, learning, and support required before implementation.
- **OPS-06.** This statement supplements existing financial and purchasing controls in response to AI-enabled impersonation and fabricated material. It does not replace the broader acquisition and provider-lifecycle requirements applying to AI systems.
- **OPS-07.** This statement applies to in-camera sessions whether board members and staff meet online, in hybrid form, or in person. It does not create an unrestricted right to prevent every necessary recording. Administrative procedures should address when recording is necessary, how objections are considered, what alternatives are reasonably available, and how resulting records are managed.
- **OPS-08.** This statement applies the general transparency principle to official administrative work. Administrative procedures may distinguish material AI assistance from routine spelling, formatting, transcription, or similar assistance that does not materially shape the content or professional judgement communicated.
- **OPS-09.** This statement does not require continuous monitoring of every public-facing system. It prevents the division from implying that urgent disclosures will receive an immediate response when no such response arrangement exists.
- **OPS-10 through OPS-13.** OPS-10 and OPS-11 apply existing restricted-use and human-responsibility requirements to operational analytics, physical systems, and administrative agents. OPS-12 and OPS-13 distinguish continuity planning from evaluation and governance reporting. They use established division processes rather than requiring separate systems for AI.

APPENDIX B

Board AI Governance Self-Check

An instrument designed to support structured conversation about readiness, safeguards, and next actions

> **Use the self-check to move from policy to evidence and action**
>
> Part One confirms that board policy includes firm safeguards related to AI use. Part Two examines how consistently governance practices are implemented. The final pages are intended to help the board identify evidence gaps, assign a limited number of priorities, and schedule follow-up. The exercise does not produce an overall score but rather is intended to identify gaps in policy and practices that require further consideration.

Purpose

This self-check supports a structured conversation among the board, the Director of Education, and the division administrators responsible for implementation. It is intended to help participants identify safeguards that require attention, practices that are established or still developing, evidence that is missing, and a manageable set of next actions.

The self-check is not an audit, certification, legal opinion, compliance determination, or comparison among school divisions. It does not produce a validated measure of organizational performance. The response levels describe stages of implementation. They do not represent points to be combined into an overall score.

The instrument was informed by the iterative Govern, Map, Measure, and Manage functions of the NIST AI Risk Management Framework and by the domain-based maturity approach of the CoSN and Council of the Great City Schools K-12 Generative AI Maturity Tool. It has been independently adapted to Saskatchewan board governance and to the legal, policy, and safeguard positions in this Advisory. Note that the instrument is not a NIST or CoSN assessment and does not indicate compliance with either framework (National Institute of Standards and Technology, 2023; Consortium for School Networking & Council of the Great City Schools, 2024).

> **Important**
>
> A strong result in one area does not offset an unconfirmed safeguard or a serious gap elsewhere. The quality of the evidence, discussion, and follow-through matters more than the number of responses at any level.

Using the self-check

PARTICIPANTS AND PREPARATION

Who should participate

The board and Director of Education should determine who can provide the evidence and perspectives needed for the discussion. Depending on the domains being reviewed, participants may include board members and administrators responsible for education, Indigenous education, privacy, information technology and security, human resources, finance, communications, student services, accessibility, procurement, and records. Student, staff, family, rights-holder, bargaining-agent, or community perspectives should be included when their participation is needed to understand an issue or plan an appropriate response.

The exercise may be completed in one extended session or over several meetings. It is usually more useful when administration assembles available evidence and proposes preliminary responses before the board discussion. Participants may also complete an initial individual review to reveal differences in knowledge or perspective.

Before beginning

Record the scope and date of the review. Identify which AI uses, functions, or proposed uses are being considered. Administration should assemble relevant evidence where available, such as policy and administrative procedures, approved-use inventories, privacy or security reviews, contracts, incident and complaint information, professional-learning plans, accessibility findings, engagement records, budget information, and reports to the board.

Do not infer that a practice exists because participants expect that it probably occurs. Select Not enough information when evidence is unavailable or participants cannot determine the current condition. Record a range of responses and the reason for disagreement when the group cannot reach a supported conclusion.

The suggested evidence listed for each domain is illustrative rather than exhaustive. It identifies possible sources for the discussion and does not mean that every listed item is required in every division.

Recommended process

1. Define the scope: Record the date, the functions or proposed uses being reviewed, and the people participating.
2. Prepare the evidence: Administration should assemble available documents and propose preliminary responses; any gaps in the evidence should not be filled by making assumptions.
3. Complete Part One first: Any safeguard needing action or evidence is recorded before the maturity discussion begins.
4. Discuss each maturity domain: Select the response supported by current evidence. Record a range and provide reasoning when the group does not agree.
5. Identify a limited number of priorities: Distinguish board policy or resource decisions from administrative, professional, engagement, or provincial responsibilities.
6. Schedule follow-up: Record who will act, what evidence will demonstrate progress, and when the matter will return to the board.

Part One: Required safeguards

CONFIRM FIRM BOUNDARIES BEFORE CONSIDERING MATURITY

These safeguards express firm boundaries in the Advisory. They are considered separately because a safeguard should not be converted into a maturity score or offset by stronger performance elsewhere.

For each safeguard, select one status:

Code	Status	Meaning
C	**Confirmed**	Current policy, procedure, practice, and available evidence support the safeguard.
A	**Action required**	The safeguard is absent, incomplete, inconsistent, or not supported by adequate evidence.
NI	**Not enough information**	Participants cannot determine whether the safeguard is in place. Identify the evidence required and who will provide it.
NU	**No current or proposed use**	Use only where the safeguard depends on a particular system or authorized use that is not occurring or is under consideration. This option is not available for conduct prohibitions or standing response responsibilities.

Complete this check first

An Action required or Not enough information response should create a specific follow-up. Select No current or proposed use only for a safeguard tied to a particular system or authorized use. It is not a permanent exemption, and it is not available for conduct prohibitions or standing response responsibilities.

Required safeguards check

MARK ONE STATUS FOR EACH SAFEGUARD

Safeguard	C	A	NI	NU
S1. Indigenous authority and protected knowledge: AI is not used to teach an Indigenous language or to generate, reproduce, or prepare material involving ceremony, ceremonial protocol, or sacred knowledge without the direction or authorization required from the relevant rights holder or authority.	☐	☐	☐	☐
S2. AI-detection results: An AI-detection result is not treated as proof or used as the sole basis for an allegation, finding, or decision.	☐	☐	☐	☐
S3. Fair process for academic-integrity concerns: A student whose work raises an academic-integrity concern has a meaningful opportunity to respond through a fair process.	☐	☐	☐	☐
S4. Teacher responsibility for assessment and evaluation: No automated system makes a summative evaluation or final decision about student achievement without review and judgement by the responsible teacher.	☐	☐	☐	☐
S5. Personal information and approved AI uses: Personal information is entered into an AI system only through a system and use that administration has assessed and approved for that purpose.	☐	☐	☐	☐
S6. Pilots and trials: An AI pilot, trial, free tool, or newly enabled feature does not bypass applicable privacy, security, accessibility, restricted-use, procurement, contractual, or other division requirements.	☐	☐	☐	☐
S7. Intimate or sexualized synthetic depictions: Division policy prohibits the school-connected creation, solicitation, or distribution of AI-generated or AI-altered intimate or sexualized depictions.	☐	☐	☐	
S8. Support and response for harmful synthetic content: Division procedures establish support and response processes for people affected by prohibited intimate, sexualized, or other harmful synthetic content.	☐	☐	☐	
S9. Harmful AI-assisted impersonation: Division policy prohibits AI-assisted impersonation intended or reasonably likely to deceive, harass, exploit, defraud, or cause harm.	☐	☐	☐	
S10. Consequential automated decisions: No automated system makes a final decision that significantly affects a student, employee, applicant, or other person. An identifiable person with authority to reach a different conclusion reviews the relevant information and remains responsible for the decision.	☐	☐	☐	☐
S11. Behavioural and emotional inference: Automated inferences about emotion, attention, honesty, engagement, or intent are not used as a basis for assessment, discipline, employment decisions, or routine monitoring.	☐	☐	☐	

C = Confirmed A = Action required NI = Not enough information NU = No current or proposed use

A shaded cell means that the NU response is not available for that safeguard.

SAFEGUARDS REQUIRING ACTION

INFORMATION OR EVIDENCE REQUIRED

LEAD ROLE AND RETURN DATE

Part Two: Governance maturity

SELECT THE CURRENT CONDITION SUPPORTED BY EVIDENCE

For each statement, select the response that best reflects the division's current condition. These responses describe implementation and are not points to be totalled.

1 | Not yet in place	Available information indicates that the practice has not been defined or implemented.
2 | Developing	Work is underway, but the practice is partial, informal, unevenly implemented, or not yet supported by sufficient resources or evidence of implementation.
3 | Substantially in place	The practice is defined and generally implemented, but some gaps, inconsistencies, or evidence limitations remain.
4 | Established and reviewed	The practice is defined, implemented consistently, supported by evidence, and reviewed or improved at planned intervals.
NI | Not enough information	Participants cannot determine whether or to what extent the practice is in place because the necessary information or evidence is unavailable. Note: Identify the follow-up information required.
NA | Not applicable	The statement genuinely does not apply to the division's present circumstances. Note: Record the reason. Do not use NA merely because a practice has not begun.

If the group disagrees

Record the range of supported responses and the reason for the difference. Do not average different perspectives into a single number. The disagreement may itself reveal inconsistent implementation, incomplete evidence, or unclear communication.

Domain 1: Board direction and accountable governance

MARK ONE RESPONSE FOR EACH STATEMENT

Use current evidence. Mark NI when the group cannot determine the condition, and record why. Mark NA only when the statement genuinely does not apply.

Maturity statement	1	2	3	4	NI	NA
1.1 The board's policy sets clear commitments, boundaries, and oversight expectations for division uses of AI.	☐	☐	☐	☐	☐	☐
1.2 Policy and administrative procedure distinguish the responsibilities of the board, division administration, school leaders, and educators.	☐	☐	☐	☐	☐	☐
1.3 Responsibility for coordinating AI governance is assigned at an organizational level with sufficient authority and access to expertise, information, and senior leadership.	☐	☐	☐	☐	☐	☐
1.4 Administration uses a documented pathway for approving, escalating, suspending, and reviewing AI uses according to their purpose and risk.	☐	☐	☐	☐	☐	☐
1.5 The board receives proportionate information about significant AI uses, material changes, serious incidents, recurring concerns, costs, capacity pressures, benefits, limitations, and matters requiring board direction.	☐	☐	☐	☐	☐	☐
1.6 Board members receive continuing AI literacy support appropriate to their governance role.	☐	☐	☐	☐	☐	☐
1.7 Board members apply the division's expectations for approved systems, confidentiality, review, human responsibility, and disclosure to the board's own work.	☐	☐	☐	☐	☐	☐

SUGGESTED EVIDENCE: board policy and review schedule; administrative procedures; assigned responsibilities; approval and escalation records; board reports; board member learning records.

1 = Not yet in place 2 = Developing 3 = Substantially in place 4 = Established and reviewed NI = Not enough information NA = Not applicable

EVIDENCE OR EXAMPLES CONSIDERED

IMPORTANT DIFFERENCES OR LIMITATIONS

FOLLOW-UP, LEAD ROLE, AND TARGET DATE

Domain 2: Indigenous authority, relationships, and knowledge governance

MARK ONE RESPONSE FOR EACH STATEMENT

Use current evidence. Mark NI when the group cannot determine the condition, and record why. Mark NA only when the statement genuinely does not apply.

Maturity statement	1	2	3	4	NI	NA
2.1 The division has an established approach for recognizing matters involving Indigenous languages, knowledge, cultural practices, ceremony, sacred knowledge, or community-held information that may fall outside the division's sole authority.	☐	☐	☐	☐	☐	☐
2.2 Proposed AI uses are screened early enough to identify when direction or authorization is required from a First Nation, Métis government or community, or First Nations education authority.	☐	☐	☐	☐	☐	☐
2.3 The division works with the relevant rights holder to determine authority, permissions, and data-governance requirements where information has collective First Nations or Métis dimensions.	☐	☐	☐	☐	☐	☐
2.4 Staff who may encounter Indigenous knowledge or community-contributed material receive practical guidance on authority, permissions, approved practices, and escalation routes.	☐	☐	☐	☐	☐	☐
2.5 The division maintains relationships through which rights holders, Elders, Knowledge Keepers, and community representatives can inform continuing policy and practice.	☐	☐	☐	☐	☐	☐

SUGGESTED EVIDENCE: engagement agreements or protocols; decision and escalation pathways; permissions records; staff guidance; records of continuing relationships and review.

1 = Not yet in place 2 = Developing 3 = Substantially in place 4 = Established and reviewed NI = Not enough information NA = Not applicable

EVIDENCE OR EXAMPLES CONSIDERED

IMPORTANT DIFFERENCES OR LIMITATIONS

FOLLOW-UP, LEAD ROLE, AND TARGET DATE

Domain 3: AI literacy, professional learning, and well-being

MARK ONE RESPONSE FOR EACH STATEMENT

Use current evidence. Mark NI when the group cannot determine the condition, and record why. Mark NA only when the statement genuinely does not apply.

Maturity statement	1	2	3	4	NI	NA
3.1 The division has a coordinated plan for role-appropriate AI literacy learning and support for students, staff, division leaders, and board members.	☐	☐	☐	☐	☐	☐
3.2 AI literacy guidance addresses how AI systems produce outputs and where they may fail. It also addresses verification, privacy, responsible use, approved systems, and the continuing need for human judgement.	☐	☐	☐	☐	☐	☐
3.3 Professional learning is sustained, role-specific, current, and supported with sufficient time and resources.	☐	☐	☐	☐	☐	☐
3.4 Administrative procedures establish conditions for student access that reflect age, developmental readiness, educational purpose, supervision, privacy, and available evidence of benefit.	☐	☐	☐	☐	☐	☐
3.5 Administration assesses how AI implementation affects staff workload, professional judgement, working relationships, and well-being.	☐	☐	☐	☐	☐	☐
3.6 The division provides support in response to material effects of AI implementation on staff workload, professional judgement, working relationships, or well-being.	☐	☐	☐	☐	☐	☐
3.7 The division identifies material differences in access to AI literacy learning and support arising from connectivity, equipment, staffing, accessibility, language, or local capacity.	☐	☐	☐	☐	☐	☐
3.8 The division responds to identified differences in access to AI literacy learning and support through appropriate supports, alternatives, collaboration, or requests for provincial action.	☐	☐	☐	☐	☐	☐

SUGGESTED EVIDENCE: literacy and professional-learning plans; role-specific guidance; time and resource allocations; student-access procedures; workload or well-being reviews; access and support data.

1 = Not yet in place 2 = Developing 3 = Substantially in place 4 = Established and reviewed NI = Not enough information NA = Not applicable

EVIDENCE OR EXAMPLES CONSIDERED

IMPORTANT DIFFERENCES OR LIMITATIONS

FOLLOW-UP, LEAD ROLE, AND TARGET DATE

Domain 4: Academic integrity, assessment, and evaluation

MARK ONE RESPONSE FOR EACH STATEMENT

Use current evidence. Mark NI when the group cannot determine the condition, and record why. Mark NA only when the statement genuinely does not apply.

Maturity statement	1	2	3	4	NI	NA
4.1 Division guidance clearly distinguishes division-wide expectations from school-, course-, and task-specific requirements for AI use.	☐	☐	☐	☐	☐	☐
4.2 Students receive clear and timely information about the AI uses and disclosures required for the work they are asked to complete.	☐	☐	☐	☐	☐	☐
4.3 Academic-integrity concerns are addressed through a documented fair process based on relevant evidence, protection of confidentiality, and a meaningful opportunity for the student to respond.	☐	☐	☐	☐	☐	☐
4.4 Procedures define the teacher review required when AI contributes to individualized feedback.	☐	☐	☐	☐	☐	☐
4.5 Procedures define how students are informed when AI materially contributes to individualized feedback.	☐	☐	☐	☐	☐	☐

SUGGESTED EVIDENCE: academic-integrity policy and procedures; student and family communications; course or task guidance; concern-review records; AI-assisted feedback procedures.

1 = Not yet in place 2 = Developing 3 = Substantially in place 4 = Established and reviewed NI = Not enough information NA = Not applicable

EVIDENCE OR EXAMPLES CONSIDERED

IMPORTANT DIFFERENCES OR LIMITATIONS

FOLLOW-UP, LEAD ROLE, AND TARGET DATE

Domain 5: Privacy, security, and data governance

MARK ONE RESPONSE FOR EACH STATEMENT

Use current evidence. Mark NI when the group cannot determine the condition, and record why. Mark NA only when the statement genuinely does not apply.

Maturity statement	1	2	3	4	NI	NA
5.1 LA FOIP requirements are incorporated into the division's governance and approval of AI uses that handle personal information.	☐	☐	☐	☐	☐	☐
5.2 Administration completes and documents a privacy review before approving a new or materially changed AI use that handles personal information, with the depth of review reflecting the sensitivity, scale, people affected, and potential consequences.	☐	☐	☐	☐	☐	☐
5.3 The division maintains an approved-use record identifying the AI system, purpose, accountable administrator or unit, authorized users, information involved, provider, storage or processing arrangements, conditions, and required reassessment.	☐	☐	☐	☐	☐	☐
5.4 Written agreements governing AI service providers address authorized information use, access, storage and processing, subcontractors, model training or service improvement, safeguards, incident notification, retention, and verified deletion.	☐	☐	☐	☐	☐	☐
5.5 The division has incorporated approved AI systems into its established processes for access control, cybersecurity, incident response, records management, business continuity, and enterprise risk.	☐	☐	☐	☐	☐	☐

SUGGESTED EVIDENCE: privacy reviews; approved-use inventory; information-flow records; agreements; security and access-control records; retention schedules; incident and continuity plans.

1 = Not yet in place 2 = Developing 3 = Substantially in place 4 = Established and reviewed NI = Not enough information NA = Not applicable

EVIDENCE OR EXAMPLES CONSIDERED

IMPORTANT DIFFERENCES OR LIMITATIONS

FOLLOW-UP, LEAD ROLE, AND TARGET DATE

Domain 6: AI systems, procurement, and intellectual property

MARK ONE RESPONSE FOR EACH STATEMENT

Use current evidence. Mark NI when the group cannot determine the condition, and record why. Mark NA only when the statement genuinely does not apply.

Maturity statement	1	2	3	4	NI	NA
6.1 The division uses one accountable pathway for acquiring, piloting, implementing, renewing, and retiring AI systems, including free tools, trials, and AI-enabled features added to existing products.	☐	☐	☐	☐	☐	☐
6.2 Before approving a proposed AI system or use, administration documents its purpose, expected benefit, implementation requirements, total cost, and applicable legal, policy, and contractual requirements.	☐	☐	☐	☐	☐	☐
6.3 Each approval identifies the authorized system or service version, purpose, users, information categories, accountable administrator or unit, conditions of use, and required reassessment.	☐	☐	☐	☐	☐	☐
6.4 Provider agreements address service and support, accessibility, permitted uses of submitted and generated material, material changes, pricing, portability, transition, termination, and continuing obligations.	☐	☐	☐	☐	☐	☐
6.5 Staff have practical guidance concerning when copyrighted, licensed, student-created, staff-created, confidential, or community-contributed material may be submitted to an AI system and who may authorize its submission or reuse.	☐	☐	☐	☐	☐	☐

SUGGESTED EVIDENCE: acquisition and approval records; documented evaluations; pilot plans; total-cost analysis; contract clauses; renewal and exit plans; intellectual-property guidance.

1 = Not yet in place 2 = Developing 3 = Substantially in place 4 = Established and reviewed NI = Not enough information NA = Not applicable

EVIDENCE OR EXAMPLES CONSIDERED

IMPORTANT DIFFERENCES OR LIMITATIONS

FOLLOW-UP, LEAD ROLE, AND TARGET DATE

Domain 7: Prohibited and restricted uses

MARK ONE RESPONSE FOR EACH STATEMENT

Use current evidence. Mark NI when the group cannot determine the condition, and record why. Mark NA only when the statement genuinely does not apply.

Maturity statement	1	2	3	4	NI	NA
7.1 Division policy and administrative guidance clearly distinguish prohibited conduct from restricted uses that require specific approval and safeguards.	☐	☐	☐	☐	☐	☐
7.2 Administration applies a documented approval process to proposed higher-risk uses, with safeguards proportionate to the system's functions, the people affected, and the consequences of failure.	☐	☐	☐	☐	☐	☐
7.3 Staff and students receive practical information about prohibited conduct, restricted uses, approval requirements, and routes for questions or reporting concerns.	☐	☐	☐	☐	☐	☐
7.4 For each approved higher-risk use, administration documents how required review, monitoring, human oversight, logging, escalation, and service continuity will operate.	☐	☐	☐	☐	☐	☐
7.5 Response procedures address support for people affected by harmful AI-generated content or system behaviour, preservation of relevant information, notification and escalation, and reassessment or suspension of the approved use.	☐	☐	☐	☐	☐	☐

SUGGESTED EVIDENCE: prohibited-use policy; restricted-use criteria and approvals; communications and training; system-specific safeguard and continuity plans; system logs and monitoring plans; incident, support, escalation, and suspension records.

1 = Not yet in place 2 = Developing 3 = Substantially in place 4 = Established and reviewed NI = Not enough information NA = Not applicable

EVIDENCE OR EXAMPLES CONSIDERED

IMPORTANT DIFFERENCES OR LIMITATIONS

FOLLOW-UP, LEAD ROLE, AND TARGET DATE

Domain 8: Equity, transparency, and community engagement

MARK ONE RESPONSE FOR EACH STATEMENT

Use current evidence. Mark NI when the group cannot determine the condition, and record why. Mark NA only when the statement genuinely does not apply.

Maturity statement	1	2	3	4	NI	NA
8.1 Administration monitors material differences in access, service levels, benefits, costs, workloads, risks, and burdens associated with AI across schools and affected groups.	☐	☐	☐	☐	☐	☐
8.2 Identified differences in connectivity, equipment, staffing, accessibility, language, or local capacity lead to appropriate supports, alternatives, collaboration, or requests for provincial action.	☐	☐	☐	☐	☐	☐
8.3 The division maintains accessible public information about its AI policy, approved student-facing systems and purposes, the general categories of information handled, and user expectations.	☐	☐	☐	☐	☐	☐
8.4 Division procedures define when people are informed that AI materially contributed to content, recommendations, or decisions affecting them, including when disclosure supports trust, accountability, or an opportunity to question the result.	☐	☐	☐	☐	☐	☐
8.5 Students, staff, families, and community members have accessible routes to ask questions, correct information, raise concerns, make complaints, and request human review.	☐	☐	☐	☐	☐	☐
8.6 Affected groups, including students, are engaged early enough to influence consequential AI decisions.	☐	☐	☐	☐	☐	☐
8.7 The division communicates what it heard through engagement and how the input affected the decision.	☐	☐	☐	☐	☐	☐

SUGGESTED EVIDENCE: access and impact information; accessibility reviews; alternatives and supports; public notices and approved-system information; correction and complaint routes; engagement records.

1 = Not yet in place 2 = Developing 3 = Substantially in place 4 = Established and reviewed NI = Not enough information NA = Not applicable

EVIDENCE OR EXAMPLES CONSIDERED

IMPORTANT DIFFERENCES OR LIMITATIONS

FOLLOW-UP, LEAD ROLE, AND TARGET DATE

Domain 9: Division administration and business functions

MARK ONE RESPONSE FOR EACH STATEMENT

Use current evidence. Mark NI when the group cannot determine the condition, and record why. Mark NA only when the statement genuinely does not apply.

Maturity statement	1	2	3	4	NI	NA
9.1 Administration maintains current knowledge of AI systems and AI-enabled features used in human resources, finance, purchasing, meetings, records, communications, facilities, transportation, and technology operations.	☐	☐	☐	☐	☐	☐
9.2 A proposed workplace use that may materially affect employees is reviewed for its defined purpose, expected benefit, reliability, possible bias or barriers, privacy, accessibility, proportionality, and less intrusive alternatives.	☐	☐	☐	☐	☐	☐
9.3 Affected staff and bargaining agents are involved as appropriate before implementation of a workplace AI use that may materially affect duties, workload, professional judgement, monitoring, or employment.	☐	☐	☐	☐	☐	☐
9.4 Financial and purchasing processes independently verify material or unusual requests and preserve delegated authority, segregation of duties, documentation, and auditability.	☐	☐	☐	☐	☐	☐
9.5 Administrative procedures govern AI recording, transcription, summarization, retention, participant notice, official records, and use in confidential or in-camera meetings.	☐	☐	☐	☐	☐	☐
9.6 Official communications produced with material AI assistance are reviewed for accuracy, accessibility, and confidentiality before release.	☐	☐	☐	☐	☐	☐
9.7 Public-facing AI services are reviewed for safety and provide an effective route to a person when the service cannot respond appropriately or a user needs assistance.	☐	☐	☐	☐	☐	☐

SUGGESTED EVIDENCE: system inventories; workplace impact reviews; engagement records; financial controls; meeting and record procedures; communications reviews; public-service escalation procedures.

1 = Not yet in place 2 = Developing 3 = Substantially in place 4 = Established and reviewed NI = Not enough information NA = Not applicable

EVIDENCE OR EXAMPLES CONSIDERED

IMPORTANT DIFFERENCES OR LIMITATIONS

FOLLOW-UP, LEAD ROLE, AND TARGET DATE

Domain 10: Evidence, review, and continuous improvement

MARK ONE RESPONSE FOR EACH STATEMENT

Use current evidence. Mark NI when the group cannot determine the condition, and record why. Mark NA only when the statement genuinely does not apply.

Maturity statement	1	2	3	4	NI	NA
10.1 The board reviews its AI policy at a defined interval and when a material change in law, provincial direction, technology, evidence, local practice, or risk requires earlier attention.	☐	☐	☐	☐	☐	☐
10.2 Administration reassesses an approved AI use when its purpose, information practices, model, integrations, provider, subcontractors, capabilities, or other relevant conditions materially change.	☐	☐	☐	☐	☐	☐
10.3 For each higher-impact use, administration has a documented monitoring process that tests whether the use continues to serve its approved purpose and identifies material incidents, harms, burdens, costs, or provider dependencies.	☐	☐	☐	☐	☐	☐
10.4 Administration can provide evidence showing whether board policy commitments and required safeguards are being implemented.	☐	☐	☐	☐	☐	☐
10.5 The board identifies a limited set of priority actions, responsible parties, required resources, and target dates after reviewing the self-check results.	☐	☐	☐	☐	☐	☐
10.6 Progress on identified priorities is reported through established governance processes and considered during the next self-check.	☐	☐	☐	☐	☐	☐

SUGGESTED EVIDENCE: policy and reassessment schedules; monitoring and incident records; implementation evidence; action plans; resource decisions; progress reports.

1 = Not yet in place 2 = Developing 3 = Substantially in place 4 = Established and reviewed NI = Not enough information NA = Not applicable

EVIDENCE OR EXAMPLES CONSIDERED

IMPORTANT DIFFERENCES OR LIMITATIONS

FOLLOW-UP, LEAD ROLE, AND TARGET DATE

Recording and interpreting the results

BUILD A RESPONSE PROFILE, NOT A SCORE

For each domain, count the statements marked at each maturity level and record the result as the domain's response profile. Also record how many statements are marked NI or NA. Do not add the ratings together, calculate an average, or produce an overall score.

1. Required safeguards: A safeguard marked "Action required" needs attention regardless of the other results. A safeguard marked "Not enough information" requires a specific evidence request. A safeguard marked "No current or proposed use" must be considered before any such use begins.
2. Information gaps: NI responses may indicate that responsibilities, documentation, monitoring, or reporting are unclear. NI is not a low maturity rating, but it should normally lead to an information-gathering action.
3. Practices needing development: Level 1 and Level 2 responses identify practices that are absent, informal, inconsistent, or still being developed. Priority depends on the significance of the issue, not just on the number of lower responses.
4. Practices to sustain: Level 3 and Level 4 responses identify practices that are substantially in place or established and reviewed. Participants should confirm that evidence supports the response and identify any remaining limitations.
5. Differences in perspective: A substantial difference among participants may reveal inconsistent implementation, incomplete evidence, or unclear communication. Record and investigate the difference.
6. Next actions: Select a manageable number of actions based on legal and policy obligations, the significance and urgency of the risk, the people affected, and the division's capacity.

What the results mean

A strong response profile shows where governance practices are established and supported by evidence. It does not certify compliance or remove the need to address an unconfirmed safeguard, information gap, incident, or consequential risk.

Worked example of a single domain

L1	L2	L3	L4	NI	NA
1	2	2	0	1	0

This profile means that for the domain in question, one practice has not begun, two are developing, two are substantially in place, and one information gap requires follow-up. The numbers should not be used to calculate an average. It is also important to identify which gaps are most consequential rather than assuming that every lower response has equal importance.

Domain response profile

COUNT RESPONSES WITHOUT CALCULATING AN AVERAGE

For each domain, count the number of statements marked at each level. Record NI and NA separately and identify the principal priority or evidence gap.

Domain	L1	L2	L3	L4	NI	NA	Principal priority or evidence gap
1. Board direction and accountable governance							
2. Indigenous authority, relationships, and knowledge governance							
3. AI literacy, professional learning, and well-being							
4. Academic integrity, assessment, and evaluation							
5. Privacy, security, and data governance							
6. AI systems, procurement, and intellectual property							
7. Prohibited and restricted uses							
8. Equity, transparency, and community engagement							
9. Division administration and business functions							
10. Evidence, review, and continuous improvement							

L1 = Not yet in place L2 = Developing L3 = Substantially in place L4 = Established and reviewed NI = Not enough information NA = Not applicable

Board discussion and action planning

TURN FINDINGS INTO A LIMITED SET OF RESPONSIBILITIES AND NEXT STEPS

Use the completed safeguard check and domain response profiles to discuss the following questions.

1. Which safeguards, legal obligations, rights, safety concerns, privacy or security matters, or Indigenous-authority questions require prompt attention?
2. Which NI responses require administration to provide evidence or clarification?
3. Which priorities require board policy, board direction, or a resource decision?
4. Which priorities relate to administrative procedure, system approval, professional learning, or operational practice?
5. Which matters require engagement with rights holders, affected staff, bargaining agents, students, families, or communities?
6. Which gaps reflect resource, connectivity, staffing, accessibility, or local-capacity constraints?
7. Which matters require collaboration with other divisions, the SSBA, or provincial authorities?
8. What limited set of actions will be undertaken first, who is responsible, what resources are required, and when will progress return to the board?

Priority action record

Priority action	Reason for priority	Lead role	Resources or engagement	Target date	Evidence of completion

Review record

Division	
Scope of review	
Date or meeting sequence	
Participants and roles	
Evidence package prepared by	
Date of progress report to the board	
Date of next self-check	

Appendix C: Understanding AI: Plain-Language Concepts

AI has introduced many new terms into education. Much of this language comes from industry and research and is used inconsistently in product descriptions, vendor materials, media coverage, and public discussion. Shared definitions matter because the wording of a policy determines which systems, practices, and decisions it covers. This appendix provides plain-language definitions of the principal AI-related terms used in this Advisory. No technical background is assumed. The definitions describe what systems and practices do rather than relying on product names, which may change quickly. A single product may fit more than one category, and its features may change after approval.

Artificial intelligence (AI): A broad term for machine-based systems that analyze inputs and generate outputs such as predictions, content, recommendations, or decisions. AI systems vary in how independently they operate and whether they adapt after deployment (OECD, 2024). AI includes technologies built into everyday software without necessarily being labelled as AI. Filtering suspicious email, recommending the next word in a sentence, generating an image, or ranking search results may all involve AI.

Generative AI: AI that produces content, such as text, images, audio, video, or computer code, in response to instructions or examples. These systems learn patterns from large collections of existing material and use those patterns to generate responses. Generative features may be built into chatbots, search engines, word processors, learning management systems, and other commonly used software.

Large language model (LLM): An AI model that processes or generates language. LLMs learn statistical patterns from very large collections of text and generate responses by predicting likely sequences of language. They do not independently verify every claim they produce. An LLM can therefore produce writing that is clear, confident, and inaccurate. The person using the output remains responsible for checking it.

Chatbot or conversational AI: A system that interacts with a user through written or spoken conversation. Some chatbots follow predetermined scripts, while others use generative AI to produce open-ended responses. Some can also access tools, retrieve records, or initiate actions. A conversational interface does not by itself indicate that a system is accurate, private, age-appropriate, or suitable for a particular purpose.

Hallucination: The term commonly used when generative AI produces information that is fabricated, unsupported, or inconsistent with source material but presents it plausibly. Examples include invented citations, details that do not appear in an uploaded document, and plausible-sounding factual errors. Because these errors may be difficult to recognize from the wording alone, important claims must be checked against reliable sources.

Deepfake: An image, audio recording, or video generated or substantially altered with AI to depict a real, identifiable person in a way that is not an authentic record. A deepfake may alter the person's appearance, voice, words, actions, or circumstances. Such media may be created for benign, disclosed, deceptive, abusive, or fraudulent purposes, so the appropriate policy response depends on the context and likely effect. Section 6.5 recommends treating school-connected AI-generated or AI-

altered intimate or sexualized depictions and harmful impersonation with the seriousness of abuse rather than ordinary technology misuse. Section 6.7 addresses fabricated voices, images, video, and records used in fraud.

Prompt: The instruction or request a person gives a generative AI system, together with any material attached to it, such as a document, spreadsheet, image, or recording. The wording and context of a prompt affect the output. A prompt may also transmit information to the provider operating the tool. How that information is stored and used depends on the product, account type, settings, and applicable agreement. Personal information, including a student's name or work, should be entered only through an approved use of an approved tool. Section 6.3 addresses the related privacy and data-governance requirements.

Training data: The material used to develop or adjust an AI model. It may include text, images, audio, video, computer code, and other information from public, licensed, purchased, or specially prepared sources. Training data affect what a system does well, where it has gaps, and which perspectives or biases may appear in its output. Languages, cultures, and communities are not represented equally in digitized material. Some material may also have been collected without the knowledge or permission of its creators. Sections 5 and 6.4 address authority, permission, and intellectual-property considerations. Depending on the applicable agreement and settings, a provider may also use information entered by users to improve its systems.

AI agent: An AI system that can plan and carry out a series of actions toward a goal with some degree of independence. An agent may search for information, work with files, complete forms, send messages, or operate other software. The distinction from a chatbot is the agent's capacity to act beyond generating a response, although an agent may still use a conversational interface. An agent may require access to accounts, information, or division systems. Sections 6.5 and 6.7 address restrictions on permissions, autonomy, monitoring, human authorization, and accountability.

AI companion: An AI system designed or used to sustain an ongoing conversational or relational interaction rather than simply complete a task. It may present a consistent personality, remember information from earlier conversations, offer advice or emotional support, and be available through text or voice. Companion features may appear in dedicated products, general-purpose chatbots, games, and social applications. Some systems encourage repeated or prolonged interaction, which may create risks involving manipulation, disclosure of private information, emotional dependency, or the displacement of human relationships (UNICEF, 2025). Sections 6.1 and 6.5 address student well-being and the safeguards required for conversational and companion-style systems.

AI-enabled wearable device: Glasses, watches, earbuds, pins, and other accessories that combine wearable hardware with AI-enabled functions. These functions may include recording, transcription, translation, identification, content generation, location tracking, biometric sensing, or real-time prompting. Such devices may support accessibility or other legitimate purposes, but their recording and communication functions may be difficult for people nearby to recognize. Section 6.5 addresses approval, notice, privacy, accessibility, and unauthorized recording.

Algorithmic bias: A systematic pattern in which an AI system produces less accurate, less reliable, or less fair results for some people or groups. Bias may arise from training data, labels, design choices, assumptions, performance measures, or the context in which a system is used. It can occur

without discriminatory intent and may be difficult to identify from a single output. Sections 6.2, 6.6, and 6.7 address bias in assessment, oversight, and decisions concerning students, staff, and others.

Automation bias: The tendency to defer too readily to the output of an automated system, even when it conflicts with other evidence or a person's professional judgement. This tendency may increase when people are under time pressure, lack information about how the system works, or must review large numbers of recommendations. Automation bias can cause an AI-informed decision to become an AI-made decision.

Human oversight: An arrangement in which an identifiable person remains meaningfully responsible for a decision involving AI. Meaningful oversight requires enough information, competence, time, and authority to question the output, consider other evidence, correct errors, override the system, and reach a different conclusion. Simply requiring a person to click an approval button is not sufficient. Responsibility for the final decision must remain clear.

Biometric identification and behavioural inference: Biometric identification or verification uses measurable physical or behavioural characteristics, such as facial geometry, fingerprints, voice patterns, gait, or keystroke patterns, to establish or confirm identity. Behavioural inference uses signals such as facial expression, voice, movement, or interaction patterns to make claims about a person's emotion, attention, honesty, intent, engagement, or other characteristics. Identification and inference are different practices and should not be treated as interchangeable. Section 6.5 recommends specific approval requirements for biometric identification and firm limits on the use of behavioural inference as evidence.

AI literacy: The knowledge, skills, and judgement people need to understand and engage with AI responsibly. It includes understanding in general terms how AI systems produce outputs and where they can fail; using approved tools for appropriate purposes; checking accuracy, bias, source reliability, and privacy implications; recognizing the effects of AI on people and communities; and deciding when AI should not be used. Operating a tool is only one part of AI literacy. Section 6.1 addresses AI literacy for students, staff, leaders, and board members, and the companion website provides the AI Literacy and Use Continuum, an illustrative, grade-banded governance reference.

AI detection tool: A product that estimates whether text or other material displays patterns associated with AI-generated content. Performance varies according to the tool, training data, type of material, language, and whether the material has been edited or paraphrased. A detection score does not establish who created the material or prove that a student used AI. Section 6.2 recommends that a score not be treated as proof of misconduct or used as the sole basis for an academic-integrity finding.

> **A CAUTION ABOUT DESCRIBING AI IN HUMAN TERMS**
>
> AI systems can produce outputs that resemble reasoning, understanding, or intention. Words such as "intelligence" and "learns" can suggest that a system thinks as a person does. Policy should not assign human consciousness, lived experience, professional judgement, or moral responsibility to a system. Language that presents the system as the decision-maker can hide the responsibility of the people who selected, configured, approved, and relied on it. Policy should separate what the system produced from what a person decided. For example: "The system generated a recommendation; a named person reviewed the relevant evidence and made the decision."