

Specimen — illustrative example

Riverstem Learning is a fictional organisation. This document comes from the scenario described on riskwoven.com/examples, and is not an assessment of any real entity. It is advisory output to support your own decisions: not legal advice, not a certification, and not an audit.

Written and reviewed by RiskWoven.com. Written by RiskWoven.com's AI drafting agent and checked by RiskWoven.com's independent review stage before release. The review is automated. It is not a human, legal or audit review.

AI Acceptable Use Policy

Riverstem Learning

Document type	policy	Version	1.0 Draft
Owner	Director of Learning	Customer decision role	Chief Executive Officer
Effective date	Upon customer release	Review cycle	Annual
Framework basis	RiskWoven AI practical best practices		
Jurisdiction	Ontario, Canada		
Prepared for	Riverstem Learning	Industry	Education
Organization size	26-100	Locations	Not specified

Advisory draft for accountable-owner and professional review. Validate legal, regulatory, contractual, technical and operational requirements before customer release. This is not legal advice, a certification, or a compliance determination.

1. Purpose

Riverstem Learning uses AI systems to support learner outcomes and reduce administrative burden. This policy states what uses of AI are permitted, what is restricted, and what is

prohibited. It applies to every person who uses an AI system on Riverstem's behalf or in connection with Riverstem's services.

This policy does not stand alone. It sits alongside the AI Governance and Accountability Policy, the AI Risk Management Standard, the Privacy and Personal Information Policy, the Information Security Policy, the Data Classification and Handling Standard, and the Third-Party Risk Management Policy. Where those instruments set more specific requirements, those requirements govern.

2. Scope

This policy applies to:

- All Riverstem employees, contractors and volunteers;
- Partner school staff who access Riverstem AI systems;
- Any third party acting on Riverstem's behalf when using an AI system that processes Riverstem information.

It covers every AI system used in connection with Riverstem's services, whether operated by Riverstem or by a third-party provider under a Riverstem agreement. It covers all information Riverstem holds, including learner identity and contact details, guardian contact details, submitted work, assessment records and safeguarding notes.

3. Definitions

AI system — any software that uses machine learning, large language models, automated reasoning or similar techniques to generate outputs, make predictions, classify information or take actions. This includes the AI tutor service and the automated marking component used by Riverstem.

Generative AI — an AI system that produces new content — text, images, audio, code or other material — in response to a prompt. The AI tutor is a generative AI system.

Approved use — a use of an AI system that appears on the AI use case inventory maintained by the Director of Learning, has been assessed under the AI Risk Management Standard, and has been reviewed by the Director of Learning.

High-impact use — a use of an AI system whose output directly affects a learner's assessment record, progression, safeguarding status or any other consequential decision

about a named individual. Automated marking of short-answer assessments is a high-impact use. Any use the Director of Learning designates as high-impact is also a high-impact use.

Human review — a qualified person reading the AI output, exercising independent judgement, and taking responsibility for the final decision before it is acted upon or communicated. Clicking through a result without reading it does not constitute human review.

4. Accountability

The **Director of Learning** owns this policy. The Director of Learning maintains the AI use case inventory, ensures training is delivered and recorded, and reviews this policy on the annual cycle set out in Section 11.

The **Chief Executive Officer** holds ultimate accountability for Riverstem's AI governance and is the review authority for this policy.

The **IT Manager** is the security owner for AI systems. Incidents involving learner information escalate immediately to the Privacy Officer and the Chief Executive Officer.

The **Privacy Officer** is accountable for ensuring AI use complies with Riverstem's obligations under applicable Ontario privacy law and the Privacy and Personal Information Policy.

Every **teacher** who uses the automated marking system is personally accountable for the grades they confirm. Confirmation is a professional judgement, not a formality.

Every **employee and contractor** is accountable for their own compliance with this policy.

5. Approved AI Systems and Uses

5.1 The AI use case inventory

The Director of Learning maintains an AI use case inventory. The inventory records every AI system Riverstem uses, the purposes for which it is used, the information categories it processes, the risk assessment outcome, and the date of last review. The inventory is a live record; it is updated whenever a new system or use case is added, changed or retired.

5.2 Currently approved systems and uses

At the date this policy takes effect, Riverstem operates two approved AI uses:

- **AI tutor service** — answers learner questions and generates practice material within the curriculum scope defined by the Director of Learning. This is an approved use for learners aged 14 to 18 enrolled through a partner school agreement.
- **Automated marking of short-answer assessments** — generates a provisional grade and feedback for teacher review. This is a high-impact use. No grade is released to a learner or recorded in the assessment system until a teacher has completed human review as required by Section 7.

No other AI system or use is approved unless it appears on the inventory.

5.3 Adding a new system or use

A person who wishes to use an AI system or use case not on the inventory must submit a request to the Director of Learning before using it. The Director of Learning assesses the request under the AI Risk Management Standard and, where the assessment supports it, adds the system or use to the inventory. Use before inventory entry is not permitted.

6. Information Restrictions

6.1 General rule

No one may enter Riverstem information into an AI system unless that system is on the AI use case inventory and the information category is explicitly listed as permitted for that system.

6.2 Safeguarding notes

Safeguarding notes are the most restricted information category Riverstem holds. Safeguarding notes must not be entered into any AI system under any circumstances. This restriction has no exceptions.

6.3 Learner information

Learner names, school, year group, submitted work and assessment history may be processed only by AI systems operating under a current education agreement that covers that information category. Processing must remain within Canada. Nothing identifying a learner may be sent to a tool outside the applicable education agreement.

6.4 Prohibition on training data use

Learner work must not be used to train any AI model, whether operated by Riverstem or by a third-party provider. Any provider agreement that would permit training on learner work is not compatible with this policy. The IT Manager and Privacy Officer must confirm this restriction is reflected in every provider contract before the system is added to the inventory.

6.5 Data classification

The Data Classification and Handling Standard sets the classification levels that apply to Riverstem information. Where that standard sets a more restrictive handling requirement than this section, the standard governs.

SPECIMEN

7. Human Review Requirements

7.1 High-impact uses require human review before action

Every output from a high-impact AI use must be reviewed by a qualified person before it is acted upon, communicated to a learner or recorded as a decision. The reviewer takes responsibility for the final decision.

7.2 Automated marking

For automated marking of short-answer assessments:

- A teacher must review and confirm every grade before it is released to the learner or entered into the assessment record.
- Any grade the model assigns below the threshold set by the Director of Learning must be read in full by the teacher, including the learner's submitted work and the model's reasoning.
- The teacher's confirmation is recorded in the assessment system against the teacher's name and the date of review.
- A teacher who disagrees with the model's grade must record their own grade and a brief reason. The teacher's grade is final.

The Director of Learning sets the marking threshold referred to above and records it in the AI use case inventory. The threshold is the organisation's to set; it is not derived from an external framework requirement.

7.3 AI tutor outputs

The AI tutor generates responses to learner questions and practice material. These outputs are not high-impact uses in themselves, but teachers remain responsible for the curriculum accuracy of material the tutor produces when that material is incorporated into formal learning activities. Teachers who use tutor-generated material in a formal context must verify it for accuracy and curriculum alignment before use.

7.4 Safeguarding escalation

The AI tutor is configured to refuse topics outside the curriculum and to escalate any learner disclosure that suggests harm to the designated safeguarding lead. This configuration is a control, not a substitute for staff judgement. Any staff member who becomes aware of a

safeguarding concern through an AI interaction must follow Riverstem's safeguarding procedure immediately, regardless of whether the system has already escalated.

8. Prohibited Uses

The following uses are prohibited for all persons covered by this policy:

8.1 Deceptive use — Using an AI system to produce content that misrepresents its origin, fabricates facts, impersonates a person, or is intended to mislead a learner, guardian, partner school or any other party.

8.2 Profiling — Using an AI system to build a profile of a learner for any purpose other than the immediate learning task for which the system is approved. Profiling for commercial targeting, behavioural prediction or any purpose outside the education agreement is prohibited.

8.3 Unapproved consequential decisions — Using an AI system to make or recommend a consequential decision about a named individual — including progression, exclusion, safeguarding status or any assessment outcome — without the human review required by Section 7.

8.4 Unapproved systems — Using any AI system not on the AI use case inventory for any Riverstem purpose.

8.5 Restricted information in unapproved systems — Entering safeguarding notes, learner personal information or any other restricted information into a system that is not approved for that information category.

8.6 Processing outside Canada — Sending learner information to any AI system that processes or stores it outside Canada.

8.7 Unlawful use — Using an AI system in a way that breaches applicable Ontario privacy law, Riverstem's obligations under its partner school agreements, or any other legal obligation.

8.8 Unsafe use — Using an AI system in a way that creates a risk of harm to a learner, including suppressing or delaying a safeguarding escalation.

A person who is uncertain whether a proposed use is prohibited must ask the Director of Learning before proceeding.

9. Verification of AI Outputs

Every person who uses an AI system output in their work is responsible for verifying that output before relying on it. Verification means checking the output against the following dimensions:

Accuracy — Is the content factually correct? Does it reflect the curriculum or subject matter accurately? AI systems can produce plausible but incorrect content.

Bias — Does the output treat learners equitably? Does it reflect assumptions about a learner's identity, background or ability that are not supported by the evidence?

Confidentiality — Does the output contain information that should not be disclosed to the recipient? Does it inadvertently reveal one learner's information to another?

Ownership — Does the output reproduce third-party material in a way that raises intellectual property concerns? Content generated by AI may incorporate protected material.

Suitability — Is the output appropriate for the age, context and purpose? Content intended for learners aged 14 to 18 must meet the standards Riverstem applies to all learner-facing material.

Verification is not optional. A person who uses an AI output without verification and that output causes harm bears responsibility for the failure to verify.

10. Reporting Obligations

10.1 What must be reported

Every person covered by this policy must report the following to the IT Manager without delay:

- A material error in an AI output that has been acted upon or communicated before it was caught;
- Actual or suspected exposure of learner information through an AI system;
- An AI output that is harmful, discriminatory, or otherwise unsafe;
- A suspected AI incident, including unexpected system behaviour, a system producing outputs outside its approved scope, or a safeguarding escalation that the system failed to make;
- A prohibited use, whether by the reporting person or by another.

10.2 How to report

Reports are made to the IT Manager. Where the incident involves learner information, the IT Manager escalates immediately to the Privacy Officer and the Chief Executive Officer. The AI Incident Response Procedure sets out the steps that follow.

10.3 No detriment

Riverstem does not penalise a person for making a good-faith report under this section. A person who conceals an incident, or who discourages another from reporting, is in breach of this policy.

11. Training and Awareness

Every person covered by this policy must complete AI acceptable use training before using any Riverstem AI system and on a cycle set by the Director of Learning. Annual completion is offered as a starting point and is not derived from an external framework requirement; the Director of Learning sets the actual cycle and records it in the training programme.

Training covers: the approved systems and their permitted uses, the information restrictions in Section 6, the human review requirements in Section 7, the prohibited uses in Section 8, and the reporting obligations in Section 10.

The Director of Learning records training completion for every covered person. Completion records are available to the CEO and the Privacy Officer on request.

A person who has not completed the required training must not use a Riverstem AI system until they have done so.

12. Exceptions

A person who believes a business need requires a departure from this policy must submit a written request to the Director of Learning before departing from the policy. The Director of Learning assesses the request and, where the assessment supports it, records a time-limited exception in the AI use case inventory.

Every exception record states: the name of the person or role it covers, the specific requirement it departs from, the business reason, the conditions that apply, and the expiry date.

The Director of Learning reviews open exceptions on a cycle set by the Director of Learning; quarterly is offered as a starting point and is not derived from an external framework requirement. An exception that has passed its expiry date without renewal is void. The number of exceptions past their expiry date is a metric reported to the CEO as part of the annual policy review.

No exception may authorise a use that is prohibited by applicable Ontario privacy law or by Riverstem's obligations under a partner school agreement.

SPECIMEN

13. Compliance and Consequences

Compliance with this policy is mandatory. The Director of Learning monitors compliance through the metrics listed in Section 14.

A breach of this policy is a conduct matter. Riverstem addresses breaches through its standard employment or contractor management processes. Serious breaches — including prohibited use of learner information, suppression of a safeguarding escalation, or deliberate concealment of an incident — may result in termination of employment or contract and referral to relevant authorities.

A partner school staff member who breaches this policy while using Riverstem systems is reported to the partner school. Riverstem may suspend that person's access pending the partner school's response.

14. Metrics

The Director of Learning tracks and reports the following metrics to the CEO at least on the annual review cycle:

- **Approved AI use cases inventoried** — the number of use cases on the AI use case inventory, updated whenever the inventory changes.
- **Workforce AI training completion** — the percentage of covered persons who have completed the required training within the required period.
- **Percentage of required human reviews completed on time** — for automated marking, the percentage of grades confirmed by a teacher before release, measured against the total grades generated in the period.
- **Number of approved exceptions past their expiry date** — exceptions that have not been renewed or closed by their expiry date.
- **Number of material control failures and overdue corrective actions** — incidents reported under Section 10 that represent a control failure, and corrective actions that have not been completed within the deadline set at the time of the incident.

The Director of Learning sets the reporting format and cadence for interim reporting. The metrics above are the minimum; the Director of Learning may add metrics that reflect emerging risk.

15. Related Documents

This policy is read alongside the following Riverstem instruments:

- AI Governance and Accountability Policy
- AI Risk Management Standard
- AI Incident Response Procedure
- Privacy and Personal Information Policy
- Information Security Policy
- Data Classification and Handling Standard
- Third-Party Risk Management Policy
- Third-Party AI Review Procedure
- Records Retention and Disposal Policy
- Business Continuity Policy

Where any of those instruments sets a more specific or more restrictive requirement than this policy, that instrument governs for the matter it covers.

16. Review and Maintenance

The Director of Learning reviews this policy annually and whenever a material change occurs — including the addition of a new AI system, a significant change to an existing system, a material incident, or a change in applicable law.

The review record states the date of review, the name of the reviewer, the changes made or the confirmation that no changes are required, and the date of the next scheduled review.

The CEO is the review authority for this policy. The reviewed policy is published to all covered persons promptly after each review cycle.

This policy version takes effect on the date recorded in the review record.

Owner review items

The drafting process identified the following in the supplied answers. Resolve each before customer release.

Type	Detail	Where
gap	The organisation's jurisdiction is recorded as Ontario, Canada, but no specific privacy statute is named in the supplied material. The policy refers to 'applicable Ontario privacy law' throughout. Riverstem's legal reviewers should confirm which instruments apply — candidates may include Ontario's education-sector privacy legislation, federal private-sector privacy law, and any sector-specific regulation — and ensure the Privacy and Personal Information Policy reflects the confirmed obligations. This document names no statute because the supplied material does not confirm which statutes apply.	Sections 6, 8, 12
gap	The supplied material does not describe Riverstem's safeguarding procedure or identify the designated safeguarding lead by role title. Section 7.4 refers to 'Riverstem's safeguarding procedure' and 'the designated safeguarding lead'. Riverstem must confirm that a safeguarding procedure exists, that the designated safeguarding lead role is formally assigned, and that the AI tutor's escalation configuration is tested against that procedure.	Section 7.4
gap	The marking threshold below which a teacher must read a submission in full is referenced in Section 7.2 but is not set in this policy. The Director of Learning must set and record this threshold in the AI use case inventory before the automated marking system is used. Until the threshold is set, the requirement defaults to full teacher review of every submission.	Section 7.2
gap	The Records Retention and Disposal Policy governs how long AI-related records — including training completion records, exception records, review records and incident records — are retained. This policy does not set retention periods because those are owned by the Records Retention and Disposal Policy. Riverstem should confirm that the Records Retention and Disposal Policy explicitly covers these record categories.	Sections 11, 12, 16
gap	The supplied material does not specify Riverstem's physical locations. The policy states that processing must remain within Canada, consistent with the customer's stated data rules. If Riverstem operates from locations outside Canada or if any staff access systems from outside Canada, the Information Security Policy and the Data Classification and Handling Standard should address those scenarios explicitly.	Section 6.3
vague	The AI Governance and Accountability Policy and the AI Risk Management Standard are referenced as sibling documents that govern related matters. This policy does not assert what those documents say, because their text was not supplied. Riverstem must confirm that the accountability structures, risk assessment criteria and inventory requirements in those documents are consistent with the commitments in this policy, particularly the use case inventory process in Section 5 and the exception process in Section 12.	Sections 5, 12, 15
gap	The Third-Party AI Review Procedure is listed as an existing document. Section 5.3 of this policy requires a new AI system to be assessed under the AI Risk Management Standard before inventory entry. Riverstem should confirm that the Third-Party AI Review Procedure is aligned with that assessment gate and that the two instruments do not create a parallel or conflicting pathway for adding third-party AI systems.	Section 5.3

Information the organization must supply

These values were not available from the intake and are required before this document can operate:

- Confirmation of which Ontario and federal privacy statutes apply to Riverstem's processing of learner information, so that the Privacy and Personal Information Policy and this policy can be reconciled with the confirmed legal obligations. Legal review is required; this document names no statute because the supplied material does not confirm applicability.
- The role title of the designated safeguarding lead, and confirmation that a safeguarding procedure exists and is referenced in staff training materials.
- The marking threshold below which a teacher must read a learner's full submission before confirming a grade. This value must be set by the Director of Learning and recorded in the AI use case inventory before the automated marking system is used in production.
- Confirmation that the Records Retention and Disposal Policy explicitly covers AI-related records: training completion records, exception records, policy review records and AI incident records.
- Confirmation that the AI Governance and Accountability Policy and the AI Risk Management Standard are consistent with the use case inventory process and exception process described in this policy, and identification of any conflicts that must be resolved.
- Confirmation that the Third-Party AI Review Procedure and the AI Risk Management Standard assessment gate described in Section 5.3 are aligned, with no parallel or conflicting pathway for adding third-party AI systems.
- The interim reporting cadence for the metrics in Section 14, set by the Director of Learning and communicated to the CEO.
- Confirmation of Riverstem's physical operating locations, so that the Information Security Policy and Data Classification and Handling Standard can address any scenario where staff access AI systems from outside Canada.

Written and reviewed by RiskWoven.com. Written by RiskWoven.com's AI drafting agent and checked by RiskWoven.com's independent review stage before release. The review is automated. It is not a human, legal or audit review.