

Specimen — illustrative example

Aldergrove Legal 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

Aldergrove Legal

Document type	policy	Version	1.0 Draft
Owner	Director of Practice Operations	Customer decision role	Managing Partner
Effective date	Upon customer release	Review cycle	Annual
Framework basis	RiskWoven AI practical best practices		
Jurisdiction	Ontario, Canada		
Prepared for	Aldergrove Legal	Industry	Legal services
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

This policy states Aldergrove Legal's position on the acceptable use of artificial intelligence systems by its lawyers, law clerks, and all other firm personnel. It sets out which uses are

permitted, which are restricted, and which are prohibited. It establishes the obligations that apply every time a member of the firm uses an AI system in connection with firm work, client matters, or firm administration.

Aldergrove Legal uses AI to support document review on discovery sets, first-draft clause generation, and summarisation of long filings. These uses carry real risk: AI systems can produce inaccurate, biased, or confidential outputs; they can expose privileged material; and they can generate citations that do not exist. This policy exists to ensure that every use of AI at the firm is deliberate, supervised, and consistent with the firm's professional obligations to its clients and to the Law Society of Ontario.

2. Scope

This policy applies to:

- all lawyers, law clerks, legal assistants, administrative staff, and any other person working for or on behalf of Aldergrove Legal, whether employed, contracted, or seconded;
- every AI system used to perform, support, or automate any task connected to firm work, client matters, or firm administration, regardless of whether the system is firm-provisioned or personally owned;
- all information processed through an AI system, including client matter files, discovery sets, trust account records, privileged communications, and employee data.

This policy applies in the Toronto office and to any remote working arrangement. It applies to AI systems accessed through firm devices and, where firm work is involved, to any other device.

This policy does not govern the internal operation of AI systems embedded invisibly in productivity tools where the user has no direct interaction with the AI component and no firm information is extracted or transmitted to an external service. Where doubt exists, the user must treat the tool as in scope and seek guidance from the IT Manager before use.

3. Definitions

AI system. Any software that uses machine learning, large language models, natural language processing, or similar techniques to generate, classify, summarise, translate, predict, or recommend outputs from inputs supplied by a user. This includes tools accessed through a browser, an API, a desktop application, or a plugin.

Generative AI. An AI system that produces new content — text, code, images, or other output — in response to a prompt. Generative AI is the category of AI system most commonly used at Aldergrove Legal for drafting, summarisation, and document review assistance.

Approved use. A specific combination of AI system, information category, and business purpose that the Director of Practice Operations has recorded in the AI use case inventory as permitted under this policy. A use is not approved simply because the system is approved; the purpose and the information category must also be approved.

High-impact use. Any use of an AI system where the output will, or is likely to, be filed with a court or tribunal, sent to a client as advice, relied on as a statement of law, or used to make a decision that materially affects a client's legal position or a person's rights. All high-impact uses require human review by the responsible lawyer before the output is acted on.

Human review. The active, substantive assessment of an AI output by a qualified person who has the knowledge and authority to accept, correct, or reject that output. Reading an output without checking it is not human review. For legal outputs, human review is performed by the responsible lawyer on the matter.

4. Policy Owner and Accountability

The **Director of Practice Operations** owns this policy. The Director is accountable for maintaining the AI use case inventory, ensuring training is delivered and recorded, reviewing exceptions, and reporting policy metrics to the Managing Partner.

The **Managing Partner** holds ultimate accountability for the firm's AI governance and is the review authority for this policy. Where a matter of privilege is engaged by an AI incident or a proposed AI use, the Managing Partner is notified without delay.

The **IT Manager** is accountable for maintaining the technical controls that enforce this policy, for receiving and triaging AI incident reports, and for escalating to the Director of Practice Operations.

The **responsible lawyer** on each matter is accountable for every AI output used in connection with that matter. Professional accountability to the client and to the Law Society of Ontario rests with the responsible lawyer and is not transferred to any AI system or to any other member of the firm.

Every member of the firm is accountable for their own compliance with this policy.

5. Approved AI Systems and the Use Case Inventory

5.1 Approved systems only. Personnel may only use AI systems that appear in the firm's AI use case inventory. Using an AI system that is not in the inventory — including consumer tools, free-tier services, browser extensions, or personal subscriptions — for any firm work is prohibited.

5.2 The AI use case inventory. The Director of Practice Operations maintains the AI use case inventory. Each entry records the AI system, the approved business purposes, the information categories that may be processed, any conditions on use, and the date the entry was last reviewed. The inventory is the authoritative record of what is permitted.

5.3 Adding a system. Any lawyer or member of staff who wishes to use an AI system not currently in the inventory must submit a request to the Director of Practice Operations before use. The Director of Practice Operations will apply the Third-Party AI Review Procedure before adding any system to the inventory. No system is used pending that review.

5.4 Approved purpose, not just approved system. An AI system being in the inventory does not make every use of it permitted. Personnel must use each system only for the purposes recorded against it. Using an approved system for a purpose outside its inventory entry requires a new request.

5.5 Current approved uses. At the date this policy takes effect, the firm's approved AI uses are:

- document review assistance on discovery sets, using the tenant-isolated service under the firm's enterprise agreement;
- first-draft clause generation, using the tenant-isolated service under the firm's enterprise agreement;
- summarisation of long filings, using the tenant-isolated service under the firm's enterprise agreement.

The Director of Practice Operations must update the inventory whenever a use is added, changed, or withdrawn.

6. Information Rules

6.1 Privileged material. Privileged client material — including client communications, matter files, discovery sets, and anything that privilege attaches to — may only be processed in the

tenant-isolated service operating under the firm's enterprise agreement. It must never be entered into a consumer AI tool, a free-tier service, or any service not covered by the firm's enterprise agreement.

6.2 Absolutely prohibited inputs. The following categories of information must not be entered into any external AI system under any circumstances:

- client identity information;
- settlement figures and negotiating positions;
- information subject to a sealing order or court-imposed confidentiality obligation;
- trust account records;
- employee personal information.

The Data Classification and Handling Standard sets out the classification of information categories in detail. Personnel must apply that standard when deciding what may be entered into an AI system.

6.3 No client material on consumer tools. Client material of any kind — privileged or otherwise — must not be entered into a consumer AI tool. This applies regardless of whether the tool appears to offer a privacy mode or claims not to retain inputs.

6.4 No personal devices. Client material must not be processed through an AI system accessed on a personal device. This is consistent with the firm's existing rule that no client material is held on personal devices.

6.5 Data residency. AI systems used to process client material must process and store that material within Canada. The IT Manager is responsible for confirming and recording data residency for each system in the inventory.

6.6 Alignment with the Privacy and Personal Information Policy. Processing of personal information through an AI system must comply with the Privacy and Personal Information Policy. Where health information, financial information, or information about opposing parties or witnesses is involved, personnel must apply heightened care and must confirm the use is in the inventory before proceeding.

7. Obligations When Using an Approved AI System

7.1 Verify every output. Personnel must not treat any AI output as accurate, complete, or reliable without verification. The obligation to verify applies to factual claims, legal propositions, citations, and summaries.

7.2 Citations must be checked. Any citation produced by an AI tool — case name, neutral citation, reporter reference, or statutory provision — must be verified against the original source before it is used in any document, filing, or communication. A citation that cannot be verified must not be used.

7.3 Check for bias and fairness. Personnel must consider whether an AI output reflects bias that could affect the quality of advice or the fairness of a decision. This is particularly important in matters involving individuals where the output will influence a legal position.

7.4 Check for confidentiality risk. Before using or sharing an AI output, personnel must consider whether the output contains, reveals, or could be used to infer confidential or privileged information.

7.5 Check for ownership. Personnel must consider whether an AI output may reproduce third-party content in a way that raises intellectual property concerns before that output is filed, published, or sent to a client.

7.6 Human review for high-impact uses. No AI output that constitutes a high-impact use may be filed, sent to a client, or relied on as a statement of law without human review by the responsible lawyer. The responsible lawyer must actively assess the output, not merely read it. The responsible lawyer records that review in the matter file.

7.7 Responsible lawyer accountability. The responsible lawyer's accountability to the client and to the Law Society of Ontario is not reduced by the use of an AI system. The responsible lawyer remains fully accountable for every output used in connection with a matter, regardless of which member of the firm generated it.

8. Prohibited Uses

The following uses are prohibited for all personnel at all times:

8.1 Deceptive use. Using an AI system to generate content intended to mislead a court, tribunal, regulator, client, opposing party, or any other person. This includes submitting AI-generated citations without verification and presenting AI output as the product of independent human analysis when it is not.

8.2 Unlawful use. Using an AI system in a way that breaches applicable law, including privacy law, intellectual property law, and the firm's professional regulatory obligations.

8.3 Unapproved consequential use. Using an AI system to make, or to directly support, a consequential decision affecting a client's legal position without the use being recorded in the inventory and without human review by the responsible lawyer.

8.4 Unsafe use. Using an AI system in a way that creates a foreseeable risk of harm to a client, an opposing party, a witness, or any other person.

8.5 Entering prohibited information. Entering any information listed in section 6.2 into any AI system, or entering privileged material into any system other than the approved tenant-isolated service.

8.6 Circumventing controls. Attempting to bypass, disable, or work around any technical or procedural control that enforces this policy, including using a personal account to access an AI system that the firm has not approved.

8.7 Unsupervised use by non-lawyers on legal outputs. Law clerks and other non-lawyer personnel must not use AI to produce a legal output — advice, a legal opinion, a statement of law — that is not reviewed by the responsible lawyer before it leaves the firm or is placed on the matter file as a final work product.

9. Exceptions

A member of the firm who believes a specific use not covered by the inventory is necessary may request an exception from the Director of Practice Operations. The Director of Practice Operations will assess the request against the firm's risk position and professional obligations before granting it.

Every exception must be:

- documented in writing, with the business justification, the information categories involved, and any conditions on use;

- time-limited, with an expiry date set by the Director of Practice Operations at the time of granting;
- recorded in the AI use case inventory as an exception entry.

An exception that has passed its expiry date is no longer valid. The Director of Practice Operations monitors the number of exceptions past their expiry date as a policy metric and reports overdue exceptions to the Managing Partner.

No exception may authorise a use that is prohibited under section 8 or that would breach the firm's professional obligations.

10. Reporting Obligations

10.1 What must be reported. Every member of the firm must report the following to the IT Manager without delay:

- an AI output that contains a material factual error that was or could have been relied on;
- an AI output that appears to have reproduced privileged, confidential, or personal information that was not in the input;
- an AI system that behaves in an unexpected or harmful way;
- a suspected or confirmed exposure of client material through an AI system;
- any use of an AI system that the person believes may have breached this policy.

10.2 Escalation. The IT Manager escalates reports to the Director of Practice Operations. Where privilege may be affected, the Director of Practice Operations notifies the Managing Partner without delay. The AI Incident Response Procedure governs the steps taken after a report is received.

10.3 No penalty for good-faith reporting. The firm will not penalise any person who reports a concern in good faith under this section, even if the concern turns out to be unfounded.

10.4 Interaction with other procedures. Where an AI incident also constitutes a privacy breach or a security incident, the Privacy Breach Response Procedure and the Security Incident Response Procedure apply alongside the AI Incident Response Procedure. The IT Manager coordinates the response.

11. Training

Every member of the firm must complete AI acceptable use training before using any approved AI system for firm work. New joiners must complete training before they are given access to any approved AI system.

The Director of Practice Operations is responsible for ensuring training is delivered, for recording completion, and for reporting the workforce training completion rate to the Managing Partner as a policy metric.

Training must cover: the requirements of this policy; the information rules in section 6; the verification obligations in section 7; the prohibited uses in section 8; and how to report a concern under section 10.

The Director of Practice Operations sets the frequency at which refresher training is required. That frequency is recorded in the AI use case inventory and reviewed as part of the annual policy review. A starting point of annual refresher training is offered here; the firm must confirm this cadence and record it formally, as it is not derived from a specific framework requirement.

12. Policy Metrics

The Director of Practice Operations tracks and reports the following metrics to the Managing Partner on a cycle set by the Managing Partner. Annual reporting is offered as a starting point; the firm must confirm this cadence, as it is not derived from a specific framework requirement.

- **Approved AI use cases inventoried.** The number of use cases recorded in the AI use case inventory, including the number added, changed, and withdrawn in the reporting period.
- **Workforce AI training completion.** The percentage of personnel who have completed required AI acceptable use training, reported by role group.
- **Percentage of required reviews completed on time.** The percentage of high-impact use human reviews that were recorded in the matter file before the output was used, as a proportion of all high-impact uses identified in the period.
- **Number of approved exceptions past their expiry date.** The count of exception entries in the inventory whose expiry date has passed and that have not been renewed or closed.
- **Number of material control failures and overdue corrective actions.** The count of incidents or audit findings that identified a material failure of a control in this policy, and

the count of corrective actions arising from those findings that are past their target resolution date.

The Director of Practice Operations escalates any metric that indicates a material deterioration in compliance to the Managing Partner outside the regular reporting cycle.

SPECIMEN

13. Relationship to Other Firm Documents

This policy operates alongside the following firm documents. Where a conflict appears, the more specific document governs for its subject matter; where genuine ambiguity remains, the Director of Practice Operations resolves it.

- **AI Governance and Accountability Policy** — sets the firm's overall AI governance framework and accountability structure. This policy implements the acceptable use requirements within that framework.
- **AI Risk Management Standard** — sets the mandatory risk assessment requirements for AI systems. The Third-Party AI Review Procedure applies that standard before a system is added to the inventory.
- **Data Classification and Handling Standard** — governs the classification of information categories referenced in section 6 of this policy.
- **Privacy and Personal Information Policy** — governs the processing of personal information, including personal information processed through AI systems.
- **Third-Party Risk Management Policy** — governs the firm's obligations when engaging external AI providers.
- **Third-Party AI Review Procedure** — the procedure applied before any new AI system is added to the inventory.
- **AI Incident Response Procedure** — the procedure followed when a report is made under section 10 of this policy.
- **Privacy Breach Response Procedure** and **Security Incident Response Procedure** — apply where an AI incident also constitutes a privacy breach or security incident.
- **Information Security Policy** and **Access Control Policy** — govern the security controls that protect firm systems, including AI systems.
- **Records Retention and Disposal Policy** — governs the retention of records created or processed using AI systems, including matter file records of human review.

14. Compliance and Consequences

Compliance with this policy is mandatory for all personnel within scope.

A breach of this policy may constitute a disciplinary matter and will be addressed under the firm's employment and conduct processes. Where a breach engages the firm's professional obligations — including obligations to the Law Society of Ontario — the Managing Partner will determine the appropriate response, which may include notification to the relevant regulatory body.

The firm will not treat a breach as a purely internal matter where it has caused or risks causing harm to a client, an opposing party, or any other person.

The AI Risk Management Standard sets out the control requirements that implement this policy. Compliance with this policy requires compliance with that standard.

15. Review and Maintenance

The Director of Practice Operations reviews this policy annually and presents any proposed changes to the Managing Partner for review. The Managing Partner is the review authority for this policy.

The review record is maintained by the Director of Practice Operations and is available to the Managing Partner on request.

The Director of Practice Operations may initiate an out-of-cycle review at any time if:

- a material AI incident occurs;
- the firm adopts a new AI system or use case that this policy does not adequately address;
- a change in the firm's professional regulatory obligations affects the requirements of this policy; or
- a change in applicable law requires amendment.

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

16. Review Record

Version	Review date	Reviewed by	Summary of changes
1.0	Date to be recorded by Director of Practice Operations	Director of Practice Operations / Managing Partner	Initial issue

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 statutes or regulatory instruments were supplied. The policy refers to 'applicable law' and 'applicable Ontario privacy law' without naming instruments. The customer's legal reviewers should confirm which of PIPEDA, Ontario's privacy legislation, or sector-specific regulation applies to the firm's processing activities, and should confirm whether the Law Society of Ontario's specific guidance on AI use imposes obligations beyond those stated in this policy. These are candidates for the customer's leg	Sections 6, 8, 14
gap	The customer's answers identify the tenant-isolated service under the firm's enterprise agreement as the approved system for privileged material, but do not name the service or the enterprise agreement. The AI use case inventory must record the specific system and agreement. This policy refers to the approved service by description only, which is appropriate for a policy instrument, but the inventory entry must be specific enough that a reviewer can confirm which service is in use.	Sections 5.5, 6.1
gap	No framework with verifiable control identifiers was supplied. The policy is grounded in RiskWoven AI practical best practices. No control identifiers have been cited because none were present in the supplied material.	Throughout
vague	The refresher training cadence and the policy metric reporting cadence are stated as organisational parameters with a starting point offered. The firm must confirm and formally record both cadences before this policy is finalised. Until they are set, the training completion metric and the review completion metric cannot be tested by a reviewer.	Sections 11, 12
gap	The existing AI Governance and Accountability Policy and AI Risk Management Standard are referenced by title. This policy is written to align with those instruments, but their text was not supplied. The customer must confirm that the accountability structure in section 4 of this policy is consistent with the AI Governance and Accountability Policy, and that the control requirements referenced in section 14 are fully addressed in the AI Risk Management Standard.	Sections 4, 13, 14
gap	The Records Retention and Disposal Policy is referenced for retention of AI-related records, including matter file records of human review. The text of that policy was not supplied. The customer must confirm that it covers records generated through AI use, including the human review records required by section 7.6 of this policy.	Section 13

Information the organization must supply

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

- The name of the specific AI system and enterprise agreement that constitutes the approved tenant-isolated service for privileged material, to be recorded in the AI use case inventory (referenced in sections 5.5 and 6.1).
- Confirmation from the firm's legal reviewers of which privacy statutes and regulatory instruments apply to the firm's processing activities in Ontario, and whether the Law Society of Ontario has issued specific AI guidance that imposes additional obligations.
- The formal refresher training cadence (annual is offered as a starting point in section 11; the firm must confirm and record this).
- The formal policy metric reporting cadence (annual is offered as a starting point in section 12; the firm must confirm and record this).
- Confirmation that the accountability structure in section 4 of this policy is consistent with the existing AI Governance and Accountability Policy.
- Confirmation that the AI Risk Management Standard fully addresses the control requirements referenced in section 14 of this policy.
- Confirmation that the Records Retention and Disposal Policy covers records generated through AI use, including matter file records of human review required by section 7.6.
- The effective date of this policy, to be inserted in the review record in section 16 by the Director of Practice Operations.
- Any additional AI systems or use cases the firm intends to approve at or after the date this policy takes effect, so that the AI use case inventory is complete before the policy is issued.

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.