CONFIDENTIAL COMPLIANCE DOCUMENT

Canada AI Compliance — Complete

Prepared exclusively for Beaumont Ridge Insurance Brokers

Canada Jurisdiction · 11 August 2026

CONTENTS

What is in this pack

Each document below is also attached to your delivery email as its own PDF and editable Word file, so you can circulate or sign any one of them on its own.

PART 1

Your AI Compliance Profile

Beaumont Ridge Insurance Brokers is a 51-250 employee commercial and personal lines broker operating across Ontario, Quebec, Alberta, and British Columbia, with offices in Toronto and Montreal. The firm deploys AI in four materially consequential ways: an automated renewal-risk scoring and quote-pricing model that operates without broker involvement on straight cases, a claims triage routing tool, AI-assisted HR and recruitment systems, and internal knowledge management via Microsoft 365 Copilot and ChatGPT. Because the firm processes health and financial data on policyholders, makes automated pricing decisions, routes claims without human review, and screens job applicants, it sits in the highest-risk tier of Canadian AI users for privacy, anti-discrimination, and employment law purposes. Compliance gaps are active and regulable now: Quebec Law 25 automated-decision notification obligations apply immediately, Ontario O. Reg. 476/24 job-posting disclosure requirements have been in force since 1 January 2026, and OPC enforcement posture on AI and privacy is intensifying.

Maximum Fine Exposure

The maximum theoretical penalty exposure for Beaumont Ridge Insurance Brokers is calculated across the three regimes that carry the largest monetary consequences. (1) Quebec Law 25: the maximum penalty is CAD $25,000,000 or 4% of worldwide turnover for the preceding fiscal year, whichever is greater. On the fixed limb, a single enforcement action for, for example, operating the automated renewal-pricing model without the required Law 25 s. 12.1 notifications, or transferring policyholder data to OpenAI without a completed Privacy Impact Assessment, could expose the firm to up to CAD $25,000,000. The turnover-based limb (4% of worldwide turnover) cannot be computed without the firm's actual revenue figures, but for a brokerage in the 51-250 employee band the 4% limb could exceed the fixed maximum if annual revenues exceed CAD $625,000,000 — at typical brokerage revenue densities for this headcount, the fixed CAD $25,000,000 cap is likely to be the binding limb, though this should be confirmed with the firm's finance team. Multiple distinct violations (e.g., automated-decision notification failure, cross-border transfer without PIA, breach notification delay) are assessed separately, so aggregate exposure across concurrent violations could be a multiple of CAD $25,000,000. (2) PIPEDA: up to CAD $100,000 per knowing or reckless violation. If three distinct violations were found (e.g., consent failure on health data, safeguard failure on ChatGPT transfers, accountability failure on automated pricing), maximum exposure on this regime is 3 x CAD $100,000 = CAD $300,000. (3) Ontario ESA / O. Reg. 476/24: up to CAD $10,000 per contravention for repeated violations per posting. If the firm has published, for example, 10 non-compliant Ontario job postings since 1 January 2026, maximum administrative penalty exposure is 10 x CAD $10,000 = CAD $100,000 on a worst-case repeat-offence basis. The dominant single-regime exposure is Quebec Law 25 at up to CAD $25,000,000 per violation, making the renewal-pricing model's automated-decision compliance gap the firm's highest-priority remediation item.

AI Tools Inventory

Tool	Use	Risk Level	Notes
Automated renewal-risk scoring and quote-pricing model	Scores renewal risk and prices quotes on straight cases with no broker involved	**High**	This is the firm's most acute compliance exposure. Quebec Law 25 (s. 12.1 of the Act respecting the protection of personal information in the private sector, as amended) requires that individuals be notified when a decision based exclusively on automated processing is made about them, and grants a right to request human review and to have inaccuracies corrected.

Tool	Use	Risk Level	Notes
			The model ingests financial and potentially health data, meaning PIPEDA consent and purpose-limitation obligations apply in all provinces. Alberta and BC PIPA impose parallel accuracy and consent requirements. OPC guidance on AI and privacy requires that automated systems producing consequential outputs be explainable and auditable. Bias and proxy discrimination risk is material: financial and postal-code data can function as proxies for protected characteristics under human rights codes in all four provinces. A documented algorithmic impact assessment and audit trail for every automated quote decision are required.
Claims triage routing tool	Routes incoming claims to handlers or workflows without broker/adjuster involvement	**High**	Routing decisions materially affect claimants' access to timely service and may embed bias if training data reflects historical handler assignment patterns. Where routing amounts to a decision exclusively by automated processing that affects policyholders' rights, Quebec Law 25 notification obligations are triggered. The tool processes personal data (claim details, health information on injury claims), engaging PIPEDA s. 5 and Schedule 1 principles on purpose, consent, and safeguards. Provincial insurance regulators in Ontario (FSRA), Quebec (AMF), Alberta (AIC), and BC (BCFSA) have supervisory interest in claim-handling fairness; AI-driven triage that causes systemic delays or unequal treatment could attract regulatory scrutiny independent of privacy law. Audit logging of routing decisions and periodic bias testing are required controls.
AI-powered CRM / sales tools	Customer relationship management, sales pipeline, and client engagement	**Limited**	Processes personal contact and transactional data; PIPEDA and provincial PIPA consent and safeguard obligations apply. Risk is limited provided the tool is not used to make consequential decisions about individuals without human review. If the CRM scores clients for cross-sell or retention in a way that influences coverage offers, the risk level rises toward High and Quebec Law 25 automated-decision provisions may apply. Data-sharing terms with the CRM vendor must be reviewed for cross-border transfer compliance and sub-processor controls.
Tool	**Use**	**Risk Level**	**Notes**
Microsoft 365 Copilot	Internal knowledge management, drafting, and productivity	**Limited**	Client health and financial data held in Microsoft 365 tenancy is accessible to Copilot by default across permissioned locations. Without granular sensitivity labelling and access controls, Copilot can surface confidential policyholder information in responses to unrelated queries, breaching PIPEDA purpose-limitation and safeguard principles and professional confidentiality obligations. The firm must conduct a data-estate audit, apply Microsoft Purview

Tool	Use	Risk Level	Notes
			sensitivity labels to restrict Copilot's scope, and ensure the Microsoft Data Processing Agreement addresses Canadian data-residency and sub-processor requirements. Risk is Limited only if these controls are in place; it is High if they are not.
ChatGPT / OpenAI	Internal knowledge management, drafting assistance	**High**	Use of public or enterprise ChatGPT with policyholder health or financial data constitutes a cross-border transfer of personal information to the United States. PIPEDA requires comparable protection and contractual safeguards for cross-border transfers (OPC accountability guidance). Quebec Law 25 imposes a mandatory Privacy Impact Assessment (PIA) before any cross-border transfer and requires the transfer to be prohibited if the destination does not afford adequate protection. If staff are pasting client data into ChatGPT sessions, this is a live PIPEDA and Law 25 breach. The firm must issue an immediate acceptable-use policy prohibiting input of personal or confidential data, and, if the enterprise API is used, must complete a Quebec PIA and execute OpenAI's data processing addendum.
AI in HR / recruitment systems	HR, recruitment, and performance management	**High**	Two distinct obligations apply. First, Ontario O. Reg. 476/24 (in force 1 January 2026) requires that every publicly advertised job posting disclose whether AI is used to screen, assess, or select applicants; Beaumont Ridge has 51-250 employees and posts vacancies publicly, so the obligation applies to every Ontario posting now. Non-disclosure is an Employment Standards Act violation. Second, AI-driven performance management that influences compensation, promotion, or termination constitutes a decision about an individual engaging Quebec Law 25 automated-decision rights and PIPEDA accountability obligations. Recruitment AI trained on historical hiring data carries material bias risk under the Canadian Human Rights Act and provincial human rights codes. The firm must audit all HR AI vendors for disclosure compliance, bias testing records, and data processing terms.

Applicable Regulations

PIPEDA — Personal Information Protection and Electronic Documents Act (federal)

Canada's baseline private-sector privacy law applies to all personal information collected, used, or disclosed in the course of commercial activity, including across provinces that lack substantially similar legislation for federally regulated activities and for interprovincial transfers. Schedule 1 requires identified purpose, consent, limited collection, accuracy, safeguards, and accountability. OPC guidance issued in 2023-2026 extends these principles explicitly to AI systems: organisations deploying AI that makes consequential decisions about individuals must be able to explain those decisions, conduct privacy impact assessments, and demonstrate that automated outputs are accurate. Cross-border data transfers require contractual protections ensuring comparable safeguards. Bill C-36, tabled 15 June 2026, would amend PIPEDA to introduce stronger data-mobility and individual-rights provisions; it has not yet received Royal Assent and is not in force.

Key date: In force. Bill C-36 amendments: tabled 15 June 2026, not yet in force — position should be confirmed as legislation progresses. · Max penalty: CAD $100,000 per violation for knowing or reckless contraventions (PIPEDA s. 28). OPC can refer matters to Federal Court.

Quebec Law 25 — Act respecting the protection of personal information in the private sector (as amended by Bill 64)

Fully in force since September 2023. Key obligations for Beaumont Ridge: (1) s. 12.1 — where a decision is made exclusively by automated processing and produces legal or significant effects on a person, the organisation must notify the individual, disclose that the decision is automated, and on request provide the personal information used, the reasons and principal factors, and the right to have a human review the decision and correct inaccuracies; (2) mandatory Privacy Impact Assessment before any cross-border personal information transfer and prohibition on transfer if adequate protection is absent; (3) appointment of a Privacy Officer whose identity must be published; (4) mandatory 72-hour breach notification to the Commission d'accès à l'information (CAI) for any breach that presents a risk of serious injury; (5) French-language rights mean all privacy notices and AI disclosures served to Quebec clients must be available in French. The automated-decision rule is directly engaged by the renewal-pricing model and potentially the claims triage tool.

Key date: In force since September 2023. Automated-decision provisions (s. 12.1) in force. · Max penalty: CAD $25,000,000 or 4% of worldwide turnover for the preceding fiscal year, whichever is greater. The turnover-based limb cannot be computed without the firm's figures.

Alberta PIPA — Personal Information Protection Act (Alberta)

Substantially similar to PIPEDA for provincially regulated activity in Alberta. Requires identified purpose, meaningful consent, accuracy, and safeguards for all personal information processed in commercial activity. AI systems that collect or use personal information — including financial data on Alberta policyholders — must comply. No province-specific automated-decision provision equivalent to Quebec Law 25, but accuracy and accountability obligations (ss. 33-34) impose a duty to ensure AI outputs used in decisions about individuals are correct and that individuals can challenge inaccuracies. Cross-border transfers require contractual comparable-protection measures.

Key date: In force. Position should be confirmed for any amendments following the OPC-provincial coordination processes under Bill C-36. · Max penalty: CAD $100,000 per violation (Alberta PIPA s. 59).

BC PIPA — Personal Information Protection Act (British Columbia)

Substantially similar to PIPEDA for provincially regulated activity in BC. Applies to personal information of BC policyholders. Accuracy (s. 34) and safeguard (s. 34) obligations constrain AI systems that process personal data and produce outputs affecting individuals. No standalone automated-decision notification rule equivalent to Quebec, but the accountability and accuracy principles require organisations to be able to explain and correct AI-driven decisions. The BC Privacy Commissioner has signalled alignment with OPC's evolving AI guidance.

Key date: In force. · Max penalty: CAD $100,000 per violation (BC PIPA s. 56).

Ontario ESA / O. Reg. 476/24 — AI in Recruitment Disclosure

In force 1 January 2026. Every publicly advertised job posting by an employer of 25 or more employees must state whether AI is used to screen, assess, or select applicants, including where a third party conducts that screening on the employer's behalf. Beaumont Ridge, with 51-250 employees, posting vacancies publicly and using AI HR tools, must include the required disclosure in every Ontario job posting from 1 January 2026. Failure to disclose is a contravention of the Employment Standards Act. The obligation extends to third-party screening services the employer uses.

Key date: 1 January 2026 (in force). · Max penalty: Ontario ESA administrative penalties: up to CAD $1,000 per contravention for first offence; CAD $5,000 for second; CAD $10,000 for subsequent. Directors and officers can be personally liable.

Canadian Human Rights Act (federal) and provincial Human Rights Codes

AI systems used in underwriting, pricing, claims routing, recruitment, and performance management must not produce discriminatory outcomes on the basis of protected characteristics (race, sex, disability, age, etc.). Financial and health data processed by the renewal-pricing model and claims triage tool can function as proxies for protected grounds. The Canadian Human Rights Commission has signalled that algorithmic discrimination is within its jurisdiction. Provincial codes in Ontario (OHRC), Quebec (Charter of Human Rights and Freedoms), Alberta (AHRA), and BC (BCHRC) apply in parallel. An adverse-effect discrimination complaint based on AI outputs does not require proof of intent.

Key date: All in force. No AI-specific amendment date to note. · Max penalty: Canadian Human Rights Tribunal: compensation for lost wages, general damages up to CAD $20,000 for pain and suffering, and special compensation up to CAD $20,000 for wilful or reckless conduct, per complainant. Provincial tribunals have comparable ranges.

Provincial Insurance Regulatory Oversight (FSRA, AMF, AIC, BCFSA)

Beaumont Ridge is licensed across four provinces. Each provincial insurance regulator has supervisory authority over policyholder treatment, claims handling, and underwriting fairness. AI-driven pricing and claims triage that produces unfair, opaque, or discriminatory outcomes may attract regulatory action under insurance-specific conduct-of-business obligations independent of privacy law. FSRA (Ontario) and AMF (Quebec) have each signalled interest in AI use by regulated entities. The AMF has the most developed articulation, aligning with IOSCO and IAIS international guidance on AI governance for financial services.

Key date: Ongoing supervisory jurisdiction. No single AI-specific effective date. · Max penalty: Varies by province; AMF can impose administrative monetary penalties and suspend or revoke licences. Position should be confirmed with each regulator for current penalty scales.

OPC Guidance on AI and Privacy

The Office of the Privacy Commissioner of Canada has issued guidance (most recently updated through 2025-2026) requiring organisations using AI for consequential decisions to: conduct privacy impact assessments; ensure AI outputs are explainable and auditable; limit data inputs to those necessary for the stated purpose; maintain human oversight on high-stakes decisions; and disclose AI use to affected individuals in plain language. The OPC is conducting active AI-related investigations as of 2026 and has coordinated with provincial commissioners. While guidance is not binding legislation, OPC findings of non-compliance result in public reports and Federal Court referrals, creating reputational and legal exposure.

Key date: Guidance operative now; OPC active investigations ongoing as of August 2026. · Max penalty: OPC findings are not directly monetary, but Federal Court orders for compliance and reputational consequences are live risks. PIPEDA penalty regime (CAD $100,000) applies on referral.

Court and Tribunal AI Disclosure Obligations

Where Beaumont Ridge files documents in litigation or regulatory proceedings — including insurance arbitration, human rights tribunal submissions, or court filings — there are emerging duties in Canadian courts and tribunals to disclose where AI has been used to draft, research, or generate content in filings. Several Canadian courts have issued practice directions on this point. Failure to disclose AI-generated content in filings can constitute a breach of candour obligations and professional conduct rules applicable to counsel.

Key date: Position should be confirmed against the specific court or tribunal rules applicable to any given proceeding. · Max penalty: Sanctions range from cost orders to striking of pleadings; professional conduct consequences for counsel.

AI Acceptable Use Policy

Your formal AI policy. Sign it, circulate it to every member of staff, and review it annually or whenever your AI tool usage changes.

AI Acceptable Use Policy

Beaumont Ridge Insurance Brokers

Version: 1.0 | Effective Date: 11 August 2026 | Next Scheduled Review: 11 August 2027

Policy Owner: Chief Compliance Officer

This policy is issued under the authority of the Chief Compliance Officer and approved by the Executive Leadership Team. It is mandatory. Every member of staff must read it, sign the acknowledgement record, and comply with it from the effective date.

1. Scope — People, Tools and Activities Governed

This policy applies to every individual who works for or on behalf of Beaumont Ridge Insurance Brokers, including permanent employees, fixed-term employees, contractors, agency workers, and secondees, at all offices including Toronto and Montreal. It applies regardless of seniority, role or line of business.

The policy governs all use of artificial intelligence tools in connection with Beaumont Ridge's business activities, whether those tools are accessed on company devices, personal devices, or third-party platforms. It covers the following tools currently approved for use:

— The automated renewal-risk scoring and quote-pricing model (the Renewal Pricing Model), which scores policyholder renewal risk and generates indicative prices on straight-through cases without live broker intervention;
— The claims triage routing tool, which classifies and routes inbound claims;
— The AI-powered CRM and sales tool;
— Microsoft 365 Copilot, used across internal productivity and communications;
— ChatGPT and OpenAI-based tools, used for drafting, summarisation and knowledge management;
— AI functionality embedded in the HR and recruitment system.

This policy covers every activity in which these tools are used, including underwriting support, client communications, claims handling, financial analysis and reporting, internal knowledge management, HR processes, recruitment, and performance management. Any AI tool or AI-embedded feature not listed above is unapproved and must not be used until the approval process in section 3 has been completed.

2. Purpose and Legal Framework

Beaumont Ridge holds health information, financial information and other sensitive personal data about policyholders. It operates under licences in Ontario, Quebec, Alberta and British Columbia, and serves clients in both English and French. These facts place the business under a multi-layered legal framework that this policy operationalises.

The principal obligations are: PIPEDA, which governs private-sector personal information across federally regulated activities and in provinces without substantially similar legislation; Quebec Law 25 (Act respecting the protection of personal information in the private sector, as amended), which since September 2023 requires notification to an individual when a decision is made exclusively by automated processing and grants individuals the right to have a human re-examine that decision; BC PIPA and Alberta PIPA, which impose equivalent consent and accountability obligations in those provinces; Ontario's Employment Standards Act and O. Reg. 476/24, in force since 1 January 2026, which require that any publicly advertised job posting disclose whether AI is used to screen, assess or select applicants (this obligation applies to Beaumont Ridge as an employer of more than 25 persons); and OPC guidance on AI and privacy, which the Office of the Privacy Commissioner of Canada actively enforces. Federal AI-specific legislation: Bill C-27 (including AIDA) lapsed in January 2025 and the government has confirmed AIDA will not be reintroduced; Bill C-36, tabled on 15 June 2026, revives private-sector privacy reforms but has not yet received Royal Assent and is not yet in force. No standalone federal AI statute is currently in force.

Maximum penalties under Quebec Law 25 reach CAD $25 million or 4% of worldwide turnover for the preceding year, whichever is greater; the turnover-based limb cannot be computed here without Beaumont Ridge's confidential financial figures, but the monetary exposure is material and this policy is a primary mitigation.

3. Approved Tools and the Process for Requesting a New Tool

Only the tools listed in section 1 are approved for use. No member of staff may introduce, trial, or use any AI tool, AI-embedded product, AI plugin, or AI API connection for business purposes without prior written approval.

To request approval for a new tool, the requesting manager must submit a completed AI Tool Assessment Form to the Chief Compliance Officer. The form must identify the tool, its vendor, the data it will process, its intended use, the jurisdictions affected, and the vendor's data processing terms. The Chief Compliance Officer will assess the request against privacy law requirements (including whether a Privacy Impact Assessment is required under PIPEDA and Quebec Law 25), bias and discrimination risk, security classification, and regulatory exposure. No tool may be used until written approval is given and, where required, a Data Processing Agreement is executed with the vendor.

Approval decisions are recorded in the AI Tool Register maintained by the Compliance team. The Register is reviewed at each annual policy review and whenever a tool is modified, replaced or discontinued. Where a vendor materially changes the functionality of an approved tool, the requesting business unit must notify the Chief Compliance Officer within five business days so that the Register entry and any associated Privacy Impact Assessment can be updated.

4. Prohibited Inputs — What Must Never Be Entered into an AI Tool

The following categories of information must never be entered into any AI tool, including Microsoft 365 Copilot and ChatGPT, unless the specific tool has been approved for that data category in writing by the Chief Compliance Officer and the required contractual and technical safeguards are confirmed in place.

Personal data of customers or staff. This includes names, addresses, dates of birth, policy numbers, claim reference numbers, health or medical information, financial account details, income figures, credit information, IP addresses, and any other information from which a living individual could be identified. Health and financial data held on policyholders is particularly sensitive; entering it into a general-purpose AI tool would constitute a disclosure under PIPEDA and, in Quebec, would likely trigger breach-of-confidence obligations under Law 25.

Confidential client material. This includes quote terms, renewal schedules, claims reserves, coverage structures, and any strategic or commercial information provided by a client in the course of the brokerage relationship. Beaumont Ridge owes a duty of confidentiality as a licensed broker; AI tools whose outputs may be used to train third-party models present a direct risk of that information leaving the firm's control.

Credentials and authentication data. Passwords, API keys, MFA codes, system tokens and similar authentication items must never be entered into any AI tool under any circumstances.

Information subject to a non-disclosure agreement. Where Beaumont Ridge has signed an NDA — whether with a client, insurer, or third party — the information covered by that NDA is prohibited from AI tool input unless the NDA expressly permits it and the relevant tool is approved for that category.

Any document or data set whose entry would breach a court order, regulatory direction, or legal professional privilege. Staff who are uncertain whether a particular input is permissible must ask the Chief Compliance Officer before proceeding, not after.

5. Verification of Output Before Use and Responsibility When Output Is Wrong

AI tools produce outputs that may be incomplete, inaccurate, biased, or legally non-compliant. No AI-generated output — whether a drafted email, a summary of policy terms, an analysis, a recommendation, or a generated document — may be used externally, submitted to a regulator, or relied upon in a client-facing decision without a qualified human review.

For outputs used in financial analysis and reporting, the accountable reviewer is the senior person responsible for the relevant report. They must confirm that figures are reconciled to source data and that no AI-generated number has been passed through unchecked.

For outputs used in customer communications, the accountable reviewer is the broker or account manager in whose name the communication will be sent. Sending an AI-drafted message is treated as the sender's personal attestation that its content is accurate and compliant.

For outputs from the Renewal Pricing Model, the accountable reviewer is the underwriting or broking lead designated for that book of business. Where the model's output is used without live broker intervention (a straight-through case), the designated lead remains responsible for the parameters and override thresholds set for that model, and must review exception reports at intervals specified in the model's operating procedure.

For claims triage routing, the claims manager is responsible for reviewing the triage tool's routing logic and confirming at the point of periodic audit that the tool is routing accurately and without discriminatory effect.

Where an AI output is wrong and causes loss, regulatory exposure or client detriment, the responsibility sits with the individual who used the output without adequate verification, and secondarily with the manager who set or approved the verification standard. The fact that an AI tool produced the output is not a defence. Errors resulting from AI use must be reported under section 9 and, where they constitute a privacy breach, under Beaumont Ridge's Privacy Breach Response Procedure.

6. Where a Human Decision Is Mandatory

Certain decisions must be made or confirmed by a human being and may not be delegated entirely to an AI system. These are not advisory suggestions; they are mandatory controls.

Renewal pricing and coverage decisions: where the Renewal Pricing Model generates a price or risk score on a non-straight-through case, a licensed broker must review and confirm the output before it is communicated to the client. On straight-through cases, the model may proceed within pre-approved parameters, but a human must review exception reports and any case that falls outside normal-distribution thresholds before the position is finalised.

Claims routing decisions with coverage implications: the claims triage tool may classify and route. It may not make a coverage determination, a denial, or a settlement authorisation. Those decisions require a qualified claims handler.

All automated decisions about individuals under Quebec Law 25: where a decision is made exclusively by automated processing and produces a legal or similarly significant effect on an individual — including a policyholder or job applicant — the individual must be notified of that fact. If the individual requests human re-examination, that request must be honoured. The Compliance team maintains the notification templates and the re-examination procedure.

HR and recruitment decisions: the AI-embedded HR system may assist in shortlisting, but no offer of employment, no rejection of an application, no performance rating, and no disciplinary or dismissal decision may be taken solely on the basis of AI output. A named human decision-maker must review and sign off every such outcome. All publicly advertised job postings must include the disclosure required by O. Reg. 476/24 (in force 1 January 2026): each posting must state whether AI is used to screen, assess or select applicants for that role. This obligation applies to Beaumont Ridge because it employs more than 25 persons, and it extends to any third-party recruiter screening on Beaumont Ridge's behalf.

Legal and regulatory filings: no document submitted to a court, tribunal, regulator or licensing body may contain AI-generated content unless it has been reviewed and verified by the responsible authorised individual. Where a court, tribunal or regulatory process requires disclosure of AI use in a filing, that disclosure must be made.

7. Disclosure — Telling Customers That AI Was Used

Beaumont Ridge is required by law and by good practice to be transparent with customers about AI use in decisions that affect them.

Under Quebec Law 25, where a decision is made exclusively by automated processing, the affected individual must be informed of that fact at or before the time the decision is communicated. The notification must explain that the decision was made exclusively by automated means and must set out the person's right to request human re-examination. Template notifications in both English and French are held by the Compliance team and must be used without unauthorised modification. Beaumont Ridge serves clients in French; all mandatory AI disclosures delivered to Quebec-based clients must be provided in French as required by the Charter of the French Language.

Beyond the statutory minimum, Beaumont Ridge's standard is that clients are informed whenever AI has played a material role in a decision or recommendation that affects their coverage, their claims or their cost. Material role means that the AI output directly influenced the decision or was the primary source of the information on which the decision rested. Brokers must not represent an AI-generated recommendation as their own unaided professional opinion.

For digital customer service interactions where a chatbot or AI-powered assistant handles the initial contact, the system must identify itself as automated at the outset of the interaction. Customers must be given a clear route to reach a human broker if they request it.

Disclosures made under this section must be recorded in the AI Activity Log described in section 8.

8. Record Keeping and Retention for AI-Assisted Work

Accountability under PIPEDA, Quebec Law 25, BC PIPA and Alberta PIPA requires that Beaumont Ridge be able to demonstrate how personal information was used in AI-assisted decisions. Records must therefore be kept at both the tool level and the transaction level.

At the tool level, the Chief Compliance Officer maintains the AI Tool Register, which records for each approved tool: the tool name and vendor, the approval date, the data categories it processes, the jurisdictions in which it is used, the Privacy Impact Assessment reference, the Data Processing Agreement reference, and the date of the last review.

At the transaction level, the business unit responsible for each AI-assisted decision must retain: a record of the input provided to the tool (subject to section 4 restrictions on what may be entered); the output generated; the name of the human reviewer who verified the output; the decision taken; and, where a Quebec Law 25 automated-decision notification was issued, a copy of that notification and any subsequent request for re-examination and the outcome.

For HR and recruitment, records must include the job posting as published (with the AI disclosure required by O. Reg. 476/24), any AI-generated shortlisting outputs, the human reviewer's sign-off, and the final decision. These records must be retained for a minimum of three years after the recruitment process closes, consistent with Ontario ESA obligations and OPC guidance.

All AI-activity records are retained in accordance with Beaumont Ridge's Document Retention Schedule. Where a record contains personal information, it is subject to the same access, correction and deletion rights as any other personal information held by the firm. Records must not be destroyed in anticipation of, or during, any regulatory investigation, legal proceedings or privacy complaint.

9. Breach of This Policy — Reporting Route and Disciplinary Position

A breach of this policy is a serious matter. Breaches include: using an unapproved AI tool for business purposes; entering prohibited categories of information into any AI tool; using an AI output externally without required human verification; failing to make a mandatory disclosure to a customer; making an automated decision about an individual without the required human review or notification; and failing to record AI-assisted work as required.

Any member of staff who becomes aware of an actual or suspected breach of this policy must report it promptly to their line manager and to the Chief Compliance Officer. Where the breach also constitutes or may constitute a personal data breach, the Privacy Breach Response Procedure must be activated simultaneously. Quebec Law 25 and PIPEDA both impose notification obligations to the relevant privacy regulator where a privacy breach presents a real risk of significant harm; delay in internal reporting directly increases the risk of regulatory non-compliance.

The Chief Compliance Officer will assess each reported breach, record it in the Breach Register, determine whether regulatory notification is required, and initiate any remediation. The individual who experienced or identified the breach will receive a written acknowledgement within two business days.

Breaches of this policy will be addressed through Beaumont Ridge's disciplinary procedure. The severity of the disciplinary outcome will reflect the seriousness of the breach, any harm caused, and whether the breach was deliberate, negligent or inadvertent. Deliberate breaches — for example, wilfully entering client health data into an unapproved consumer AI tool — may constitute gross misconduct. No disciplinary action will be taken against a member of staff who reports a breach in good faith, even if they were peripherally involved.

10. Review Cycle, Policy Owner and Signature Block

This policy will be reviewed annually, with the next scheduled review on 11 August 2027. It will also be reviewed promptly whenever any of the following occurs: a new AI tool is approved for use; a material change is made to an approved tool's functionality; a relevant law or regulatory obligation changes; a significant AI-related incident or breach occurs; or the OPC, the Commission d'accès à l'information du Québec, or any other relevant regulator issues guidance or initiates an investigation that affects Beaumont Ridge's obligations.

Bill C-36, tabled 15 June 2026, has not yet received Royal Assent. When it does, the Chief Compliance Officer must assess its impact on this policy and on Beaumont Ridge's AI governance arrangements and update this policy accordingly within 60 days of commencement.

Policy Owner: Chief Compliance Officer, Beaumont Ridge Insurance Brokers.

The Policy Owner is responsible for maintaining this document, communicating updates to all staff, ensuring the AI Tool Register and Breach Register are current, and reporting on AI compliance to the Executive Leadership Team at least annually.

Approval and signature:

Name: ________________________________ Title: Chief Compliance Officer
Signature: ____________________________ Date: ________________

Name: ________________________________ Title: Chief Executive Officer
Signature: ____________________________ Date: ________________

All members of staff must sign a dated acknowledgement confirming they have read and understood this policy. Acknowledgement records are retained by Human Resources for the duration of employment plus three years.

Employee AI Guidelines — Practical Dos and Don'ts

The working-level guidance that sits under the policy: concrete examples of good and bad AI use in this business.

Introduction and Scope

These guidelines apply to every employee and contractor at Beaumont Ridge Insurance Brokers who uses an AI tool in the course of their work. They translate the company's AI Policy into concrete, day-to-day instructions. The tools in scope are: the automated renewal-risk scoring and quote-pricing model; the claims triage routing tool; the AI-powered CRM and sales tools; Microsoft 365 Copilot; ChatGPT / OpenAI; and the AI system used in HR and recruitment. Where these guidelines say 'you must' or 'you must not', that is a compliance requirement, not a suggestion. Breaches may expose Beaumont Ridge to regulatory action under PIPEDA, Quebec Law 25, BC PIPA, Alberta PIPA, Ontario ESA / O. Reg. 476/24, and OPC guidance, as well as professional licensing consequences in Ontario, Quebec, Alberta and British Columbia. Quebec Law 25 carries a maximum fine of CAD $25 million or 4% of global turnover, whichever is higher; the turnover-based limb of that calculation depends on Beaumont Ridge's own figures and must be assessed separately by management.

Worked Examples and Do / Do-Not Pairs for Each AI Use

FINANCIAL ANALYSIS AND REPORTING — The automated renewal-risk scoring and quote-pricing model ingests policyholder financial and claims data to produce a risk score and a quoted premium without broker involvement on straight cases. This model makes or materially influences decisions about pricing, which affects individual policyholders and constitutes automated decision-making under Quebec Law 25, Article 12.1, which requires that a person whose rights are affected by a decision based exclusively on automated processing be informed of that fact and be given the opportunity to have a human review it. Do: before issuing a renewal or quote produced entirely by the model on a straight case, confirm that the policyholder disclosure has been given and that the human-review pathway is documented and accessible. Do not: treat the model's output as final and dispatch it to a client without that disclosure and pathway in place, even when no broker has been involved in the case.

CUSTOMER SERVICE AND CHATBOT INTERACTIONS — AI-powered CRM tools and any chatbot channel interact directly with clients, some of whom communicate in French. Quebec's Charter of the French Language requires that clients in Quebec be able to receive services in French, including AI-generated responses. Do: configure customer-facing AI outputs to detect the client's language and respond in French where required; review AI-drafted responses before they are sent to confirm the language is correct, the facts are accurate, and no confidential information about another client has been included. Do not: allow the AI tool to send an automated response in English only to a Quebec policyholder, and do not publish AI-generated client communications without a named employee reviewing them first.

HR, RECRUITMENT, AND PERFORMANCE MANAGEMENT — Beaumont Ridge posts vacancies publicly and employs more than 25 people, placing it squarely within Ontario ESA / O. Reg. 476/24, which came into force on 1 January 2026. Every publicly advertised job posting must state whether AI is used to screen, assess or select applicants. Do: include a clear disclosure in every public job advertisement stating that AI tools are used in the recruitment process and describing what stage they operate at (for example, initial CV screening). Do not: post a vacancy, internally or publicly, that omits AI screening disclosure if an AI system is in fact used at any stage; and do not allow an AI system to make a final hiring or dismissal decision without a human decision-maker reviewing and approving it.

INTERNAL KNOWLEDGE MANAGEMENT — Microsoft 365 Copilot is used to draft internal briefings, summarise policy documents, and support broker research. Do: use Copilot to produce first drafts and summaries that a qualified broker then reviews and edits before the content is relied upon or shared. Do not: copy policyholder names, policy numbers, health information, or financial data from client files into a Copilot prompt, as this routes personal data through Microsoft's processing environment in ways that may not be covered by existing data-processing agreements.

Prompting Practice: Keeping Personal and Client Data Out of AI Inputs

Every time you type a prompt into ChatGPT, Microsoft 365 Copilot, or any other AI tool, you are submitting data to a processing environment that may not be subject to the same access controls as Beaumont Ridge's own systems. PIPEDA Principle 4.7 (Safeguards) and Quebec Law 25 both require that personal information be protected by security measures appropriate to its sensitivity. Health information and financial information — both of which Beaumont Ridge holds on policyholders — are among the most sensitive categories recognised by the OPC.

The rule is: anonymise before you prompt. If you need the AI to help you draft a coverage analysis for a commercial client, describe the scenario without including the client's name, policy number, address, revenue figures, or any information that could identify them. Write 'a manufacturing company with three premises' rather than the client's actual name and postcode. If the task genuinely cannot be completed without real identifiers, it must be handled within a system that has a signed data processing agreement covering that data type — raise this with your line manager before proceeding.

For HR tasks, do not paste a candidate's CV, performance review, or medical information into ChatGPT or any tool without a confirmed data processing agreement. Use only anonymised or synthetic examples when testing prompts or exploring how an AI tool responds to a particular type of input. Keep a brief note of any prompt that you realise, after the fact, contained personal data, and report it to the Privacy Officer so it can be assessed as a potential breach under PIPEDA and Quebec Law 25.

Checking Output for Fabrication, Bias, and Confident Errors

AI language models, including ChatGPT and the generative features within Microsoft 365 Copilot, produce text that reads as authoritative regardless of whether it is accurate. They can invent case citations, misstate statutory provisions, confuse policy wordings, and present outdated regulatory positions with complete confidence. In the insurance brokerage context, a fabricated policy condition or an incorrect premium figure passed on to a client is not merely embarrassing — it may constitute a misrepresentation with professional liability consequences.

For every piece of AI-generated content that will be used in client-facing communications, regulatory filings, or internal decisions, the reviewing employee must check: (1) every factual claim that can be independently verified, including regulatory requirements, coverage terms, and financial figures; (2) whether the output could reflect bias against a protected characteristic — for example, whether the renewal-risk model is scoring certain postal codes, occupations, or demographic proxies in ways that correlate with race, ethnic origin, or disability, which would engage the Canadian Human Rights Act and provincial human rights codes; and (3) whether the tone or content of a customer-service response could disadvantage a client who is a member of a vulnerable group.

For the automated renewal-risk scoring model specifically, the underwriting or analytics team must conduct periodic bias audits comparing outcomes across postal codes and client profiles. Any pattern suggesting systemic disadvantage to a protected group must be escalated immediately to the Chief Compliance Officer and the Privacy Officer, and must not be concealed in a regulatory filing or customer communication. The OPC has active AI investigations under way as at the date of these guidelines, and demonstrating a documented audit trail is a material part of showing good-faith compliance.

Handling AI Output That Will Be Seen by a Customer

Any communication to a policyholder, claimant, or prospective client that has been drafted or substantially shaped by an AI tool requires a named broker or customer service professional to review and approve it before it is sent. This applies whether the channel is email, the client portal, an AI-powered CRM automated message, or a chatbot response.

Before approving AI-generated customer output, the reviewing employee must confirm: the information is factually correct and consistent with the client's actual policy; the language is appropriate to the client's preferred language (French for Quebec clients, as required); no confidential information about another client has been included; and the tone does not mislead the client about the extent of their coverage or their rights in a claim. Where the output concerns a renewal quote or pricing decision produced by the automated model on a straight case, the employee must also confirm that the Quebec Law 25 automated-decision disclosure has been given and that the client has been told they may request human review.

For claims triage, the routing tool may direct a claim to a particular handler or workflow, but the initial communication to the claimant must be reviewed by a human before dispatch. A claimant who receives an AI-generated message denying or limiting their claim and who later asks how that decision was reached must be able to obtain a meaningful explanation. Employees must never tell a claimant that 'the system decided' without being able to explain the criteria the system applied, as this would fail the Quebec Law 25 transparency requirement and may breach OPC guidance on explainability in automated decisions.

When to Stop and Escalate to a Person

The following situations require the employee to stop using the AI tool and escalate to a manager, the Privacy Officer, or the Chief Compliance Officer without delay.

First, if the automated renewal-risk scoring or quote-pricing model returns an output that appears anomalous — for example, a premium that is dramatically higher or lower than a comparable manual quote, or a risk score that cannot be explained by the visible case facts — do not send it to the client. Log the case and refer it to the underwriting supervisor for manual review.

Second, if you realise that a prompt you have submitted to ChatGPT, Copilot, or any external AI tool contained personal data, a client name, a policy number, health information, or financial data, stop and report it to the Privacy Officer immediately. Under PIPEDA and Quebec Law 25, a breach of security safeguards that creates a real risk of significant harm must be reported to the OPC and, under Law 25, to the Commission d'accès à l'information, within 72 hours of the organisation becoming aware of it.

Third, if the claims triage tool routes a claim in a way that appears to discriminate against a claimant on the basis of a protected characteristic, or if a client raises a complaint that their claim was treated differently from what they expected on grounds they believe are discriminatory, escalate to the compliance team before any further automated processing occurs on that file.

Fourth, if an AI-generated customer communication has already been sent and you subsequently discover it contained an error — incorrect coverage information, a wrong premium, or a factual misstatement — notify your line manager and the compliance team at once. Do not attempt to correct it with a further AI-generated message without human oversight.

Fifth, for HR and recruitment, if the AI recruitment tool produces a shortlist or scoring that you believe may reflect bias against candidates on the basis of a protected characteristic under the Canadian Human Rights Act or applicable provincial human rights legislation, halt the process and refer the matter to HR leadership and the Chief Compliance Officer before any offers or rejections are communicated.

Quick Reference: Key Rules at a Glance

Anonymise client and candidate data before entering any prompt into ChatGPT, Copilot, or any AI tool that does not have a signed data processing agreement covering that data type. Always include the automated-decision disclosure and human-review pathway for Quebec policyholders whose renewal or quote was produced entirely by the model, as required by Quebec Law 25. Every public job posting must disclose AI use in screening, assessment or selection, as required by Ontario ESA / O. Reg. 476/24, which has been in force since 1 January 2026. Review every AI-generated customer communication before it is sent — do not allow fully automated dispatch without human approval. Check AI output for fabricated facts, incorrect policy terms, and signs of bias before relying on it. Serve Quebec clients in French, including in AI-generated communications. Escalate anomalous model outputs, personal-data prompt incidents, suspected bias, and post-send errors to the Privacy Officer or Chief Compliance Officer immediately. These rules are not exhaustive; they sit under the Beaumont Ridge AI Policy and do not replace it.

AI Tool Inventory

The live register of every AI tool in the business. Regulators, insurers and enterprise customers all ask for this first.

Tool	Vendor	Business purpose	Data it touches	Personal data (Yes/No)	Risk level	Owner	Date reviewed
Automated renewal-risk scoring and quote-pricing model	To be confirmed — record vendor name and version	Scores renewal risk and prices quotes on straight cases with no broker review. Triggers automated pricing decisions.	Policyholder name, policy history, claims data, financial information, location.	Yes	HIGH — makes decisions about individuals with no human broker involved; engages Quebec Law 25 automated-decision notification duty and PIPEDA accountability obligations.	Chief Operating Officer	11 August 2026
Claims triage routing tool	To be confirmed — record vendor name and version	Routes incoming claims to the appropriate handler or queue without broker intervention.	Policyholder name, claim details, policy data, potentially health information.	Yes	HIGH — processes sensitive health and financial data; influences claim outcomes; bias risk in routing.	Head of Claims	11 August 2026
AI-powered CRM / sales tools	To be confirmed — record vendor name and version	Supports pipeline management, lead scoring and client engagement tracking across commercial and personal lines.	Client contact details, policy interests, interaction history, financial appetite data.	Yes	MEDIUM — supports broker decisions but outputs are reviewed by staff before action.	Head of Sales	11 August 2026
Microsoft 365 Copilot	Microsoft Corporation	Drafts emails, summarises documents, generates internal reports and supports	Email content, documents, meeting transcripts, internal data — may	Yes	MEDIUM — broad data access across the Microsoft 365 tenant; staff must be trained not to	Head of IT / Data Protection Officer	11 August 2026

Tool	Vendor	Business purpose	Data it touches	Personal data (Yes/No)	Risk level	Owner	Date reviewed
		knowledge management across the firm.	include policyholder or employee personal data.		input client confidential data without authorised controls.		
ChatGPT / OpenAI	OpenAI, L.L.C.	Ad hoc drafting, research support, internal knowledge queries and customer-facing content drafting.	Whatever staff input — risk of inadvertent entry of client, policy or employee personal data.	Yes	MEDIUM-HIGH — no dedicated enterprise data-processing agreement confirmed; data residency outside Canada unless OpenAI Enterprise terms engaged.	Head of IT / Data Protection Officer	11 August 2026
AI in HR / recruitment systems	To be confirmed — record vendor name and version	Screens, assesses or selects applicants for publicly advertised vacancies; supports performance management.	Applicant CVs, contact details, assessment scores, employee performance records.	Yes	HIGH — Ontario O. Reg. 476/24 (in force 1 January 2026) requires public job postings to disclose AI use in screening, assessment or selection. Bias and discrimination risk.	Head of Human Resources	11 August 2026
[Add new tool here]							
[Add new tool here]							

DOCUMENT TITLE: AI Tool Inventory — Beaumont Ridge Insurance Brokers | Version 1.0 | Date: 11 August 2026

RISK RATING SCALE
LOW: AI assists with internal tasks; no personal data; no decisions affecting individuals.
MEDIUM: Personal data processed; human review occurs before any decision is acted upon.
MEDIUM-HIGH: Personal data processed; some automation; data residency or contractual gaps present.
HIGH: AI makes or materially influences decisions about individuals; sensitive data involved; regulatory notification or disclosure duties engaged.

REGISTER OWNER AND REVIEW CADENCE
This register is owned by the Data Protection Officer (or nominated compliance lead). It must be reviewed in full every six

months and updated within 14 days of any new AI tool being procured, any material change to an existing tool's function or data inputs, or any change in applicable law. The next scheduled review is due by 11 February 2027.

KEY LEGAL OBLIGATIONS REFLECTED IN THIS REGISTER

1. PIPEDA (federal): Principle 4.1 (accountability), Principle 4.7 (safeguards) and Principle 4.9 (individual access) apply to all tools processing personal data. OPC guidance on AI and privacy (2024–2026) requires documented accountability for automated systems.
2. Quebec Law 25: Where a decision is made exclusively by automated means, affected individuals must be notified and given the right to request human review (s. 12.1 of the Act respecting the protection of personal information in the private sector, as amended). The renewal-risk pricing model is the primary trigger. Maximum penalty: CAD $25,000,000 or 4% of worldwide turnover for the preceding fiscal year — the turnover-based limb cannot be computed without Beaumont Ridge's figures.
3. Alberta PIPA and BC PIPA: Apply to personal data collected, used or disclosed by Beaumont Ridge in those provinces. Comparable accountability and safeguard duties to PIPEDA.
4. Ontario ESA / O. Reg. 476/24 (in force 1 January 2026): Every publicly posted job vacancy must state whether AI is used to screen, assess or select applicants. Applies to Beaumont Ridge as an employer of 25 or more. Non-disclosure is a violation. Audit all current and future job postings immediately.
5. Federal AI legislation: Bill C-27 (which included the Artificial Intelligence and Data Act, AIDA) lapsed in January 2025. AIDA will not be reintroduced. Bill C-36, tabled 15 June 2026, revives the private-sector privacy reforms from C-27 and would amend PIPEDA; it is not yet in force. No standalone federal AI statute is currently in force. Compliance obligations arise from the privacy statutes listed above and OPC guidance.
6. Bias and discrimination: The renewal-risk model and HR recruitment system both carry discrimination risk under provincial human rights codes (Ontario, Quebec, Alberta, BC). Outputs must be audited for disparate impact, particularly where health or financial data correlates with protected characteristics.
7. Client confidentiality: Microsoft 365 Copilot and ChatGPT must not be used to process confidential client data unless a compliant data-processing agreement, with Canadian data residency or adequate transfer protections, is in place. Review OpenAI Enterprise terms immediately.
8. French-language service: Quebec clients must be able to interact with any customer-facing AI tool, including any chatbot or claims triage interface, in French, in compliance with the Charter of the French Language.
9. Court and tribunal filings: Where AI has been used to draft or research any document filed with a court or tribunal, disclosure of that use is required under applicable court rules. Staff must record AI involvement in all litigation-adjacent work.

ACTION FLAGS AT DATE OF ISSUE

- Confirm vendor names and versions for rows 1, 2, 3 and 6 immediately; a register with blank vendors does not satisfy PIPEDA accountability obligations.
- Issue Quebec Law 25 automated-decision notification to policyholders affected by the renewal-risk pricing model if not already done.
- Audit all Ontario job postings for O. Reg. 476/24 disclosure compliance (mandatory since 1 January 2026).
- Obtain or confirm OpenAI Enterprise data-processing agreement before further use of ChatGPT on client or employee data.
- Schedule bias audit of renewal-risk model and HR recruitment system within 90 days.

Plain-English Summary of Your Obligations

The whole regulatory picture in language a director can act on without a lawyer sitting next to them.

What Applies to Beaumont Ridge Insurance Brokers and Why

Beaumont Ridge Insurance Brokers is subject to a layered set of privacy, automated-decision, employment and sector-specific obligations because of three facts taken together: you process personal data (including health and financial information) about policyholders; you use AI to make or materially influence decisions that affect those people; and you operate across four provinces with offices in two of them.

At the federal level, PIPEDA governs all personal information you collect, use and disclose in the course of commercial activity, including everything processed by Microsoft 365 Copilot, ChatGPT/OpenAI, your CRM, your HR systems and, critically, your automated renewal-risk scoring and quote-pricing model. PIPEDA requires a lawful basis for every processing activity, limits use to the purpose for which data was collected, and demands reasonable safeguards. The Office of the Privacy Commissioner of Canada is actively investigating AI-related privacy practices and has issued guidance that directly targets automated scoring systems of the kind you operate.

Bill C-36 was tabled on 15 June 2026 and would amend PIPEDA to introduce stronger individual rights and stricter accountability rules. It is not yet in force, so it does not impose obligations today — but you should monitor its progress and design your current programme to be capable of meeting its likely requirements.

In Quebec, Law 25 adds obligations beyond PIPEDA. The provision most directly relevant to Beaumont Ridge is the requirement to notify any person who is subject to a decision made exclusively by automated processing, to tell them of their right to have a human review that decision, and to tell them the factors that influenced the outcome. Your renewal-risk scoring and quote-pricing model, which operates with no broker involved on straight cases, falls squarely within this rule. Quebec's maximum penalty is CAD $25 million or 4% of global turnover — whichever is greater. The turnover-based limb cannot be computed without your figures, but CAD $25 million represents the floor.

Alberta and British Columbia each have their own private-sector privacy statutes (Alberta PIPA and BC PIPA) that apply to employee and customer data collected in those provinces. The obligations are broadly equivalent to PIPEDA — consent, purpose limitation, safeguards, access rights — but are enforced by provincial commissioners who can investigate independently of the OPC.

In Ontario, a regulation under the Employment Standards Act that came into force on 1 January 2026 requires every publicly advertised job posting made by an employer of 25 or more people to state whether AI is used to screen, assess or select applicants. Beaumont Ridge employs 51–250 people, posts its own vacancies publicly, and uses an AI HR/recruitment system. This obligation applies to you in full.

Federal AI-specific legislation does not currently exist. AIDA lapsed in January 2025 and the government has confirmed it will not be reintroduced. Bill C-36 covers privacy reform, not a standalone AI statute. Any advice suggesting that a federal AI Act is in force would be wrong.

Your claims triage routing tool affects how policyholder claims are handled. Although it may not be the final decision-maker, routing decisions can delay or redirect claims and should be treated as consequential automated processing for both PIPEDA and Quebec Law 25 purposes. Your CRM and sales AI tools process customer data and must comply with consent and purpose-limitation rules across all four provinces.

Obligations That Carry a Deadline

Several of your obligations are already overdue or have fixed dates you must treat as hard stops.

Ontario AI disclosure in job postings — 1 January 2026 (already in force). Every publicly advertised vacancy posted on or after that date must include a clear statement that AI is used to screen, assess or select applicants, if that is the case. Your HR/recruitment AI system means this disclosure is required now. If you have posted any vacancy since 1 January 2026 without it, those postings are non-compliant today. There is no grace period.

Quebec Law 25 automated-decision notification — in force since September 2023. The obligation to notify individuals before or at the point a decision is made exclusively by automated processing, to explain their right to human review, and to identify the factors used, has been live for almost three years. If your renewal-risk scoring and quote-pricing model has been operating without these notifications, every quote issued on a straight case without broker involvement has been issued in breach of Quebec Law 25. Remediation is required immediately.

Quebec privacy impact assessments — Law 25 requires a Privacy Impact Assessment before deploying any system that collects, uses or discloses personal information for new purposes, or that involves automated decision-making. If a PIA was not completed before your scoring model, triage tool or CRM AI went live, one must be completed retrospectively and documented.

PIPEDA breach notification — if any AI system at Beaumont Ridge suffers a security incident involving personal information that creates a real risk of significant harm, you must notify the OPC and affected individuals without unreasonable delay. This is not a future obligation: it applies to every incident that has occurred or will occur.

Bill C-36 — tabled 15 June 2026, not yet in force. No compliance deadline exists yet, but parliamentary passage is anticipated. You should treat your current programme as preparatory groundwork.

No deadline currently applies under Alberta PIPA or BC PIPA beyond the general ongoing obligations that have been in force for years. Ontario's ESA AI posting rule is the only new hard date affecting you in this reporting period.

What Happens If Each Obligation Is Missed

The consequences are concrete and, in Quebec especially, severe.

If Beaumont Ridge fails to provide the required automated-decision notification under Quebec Law 25, the Commission d'accès à l'information can impose an administrative monetary penalty of up to CAD $25 million or 4% of global turnover, whichever is greater. These figures are not theoretical: the CAI has enforcement powers and has signalled active interest in insurance sector AI use. Beyond the fine, the CAI can order you to halt the automated processing — meaning your quote-pricing model could be suspended by regulatory order, disrupting your entire straight-case workflow.

If job postings in Ontario have not included the required AI disclosure since 1 January 2026, the Ministry of Labour can issue compliance orders and levy fines under the Employment Standards Act. Reputational exposure is also immediate: job applicants who later discover they were AI-screened without disclosure have grounds to complain publicly and to regulators.

If PIPEDA consent or safeguard obligations are breached — for example, if ChatGPT or Copilot processes policyholder health or financial data beyond the scope for which consent was obtained — the OPC can investigate, publish findings, and refer the matter to Federal Court for a binding order. The OPC's AI investigations unit is active. A finding against Beaumont Ridge would be published and would directly affect your regulatory standing with provincial insurance regulators across Ontario, Quebec, Alberta and BC.

If a breach notification is missed or delayed, the OPC can treat the delay itself as a separate contravention, compounding the original breach.

Alberta and BC commissioners can investigate, audit and publish findings under their respective statutes. Although their penalty powers are currently less severe than Quebec's, a provincial finding of non-compliance is reportable to insurance regulators and can affect your licence conditions.

On bias and discrimination: if your renewal-risk scoring model produces outcomes that disproportionately disadvantage individuals on the basis of a protected characteristic — age, sex, ethnicity, disability — you face liability under provincial human rights codes across all four provinces, as well as scrutiny from insurance regulators who have duties to ensure fair treatment of consumers. A single upheld human rights complaint about algorithmic discrimination can trigger a regulatory review of your entire AI scoring methodology.

What Is Already Handled by the Documents in This Compliance Pack

This plain-English summary is one document in a complete compliance pack prepared for Beaumont Ridge. The other documents in the pack cover the following areas, so they are not repeated here in full.

The AI Register documents every AI tool in scope — the renewal-risk scoring and quote-pricing model, the claims triage routing tool, the CRM/sales AI, Microsoft 365 Copilot, ChatGPT/OpenAI, and the HR/recruitment system — together with their purpose, data inputs, decision outputs and risk classification. That register is your primary evidence of governance if the OPC, CAI or any provincial regulator asks what AI you operate and how you control it.

The Privacy Impact Assessment template and completed PIAs for your highest-risk systems address the Quebec Law 25 PIA obligation and the OPC's expectation that organisations conduct privacy-by-design assessments before deploying AI that processes personal information.

The Automated Decision-Making Policy sets out the process by which Quebec Law 25 notifications are delivered to policyholders whose quotes are generated by the scoring model without broker involvement, the mechanism for requesting human review, and how the factors influencing the automated decision are communicated. This policy is the operational instrument that keeps you on the right side of the CAI.

The Data Processing and Vendor Agreement review covers the obligations owed by Microsoft (for Copilot), OpenAI (for ChatGPT), and your CRM and HR system vendors as processors of personal data on your behalf, including data-residency considerations relevant to your Montreal and Toronto offices and the health/financial data you hold.

The Bias and Fairness Testing Protocol addresses the AI bias risk you identified as a concern. It sets out how you test the renewal-risk model and triage tool for discriminatory outputs and how you document those tests.

The HR/Recruitment AI Disclosure Template provides ready-to-use language for Ontario job postings to satisfy the 1 January 2026 obligation under Ontario's regulation, and extends that disclosure standard to Alberta and BC postings as a matter of consistent practice.

The Incident Response Procedure covers PIPEDA breach notification, CAI breach notification under Quebec Law 25, and the equivalent provincial obligations in Alberta and BC.

The Three Things to Do First

Given the range of obligations and the fact that some are already overdue, Beaumont Ridge must prioritise. These three actions carry the greatest legal exposure if left undone and must be completed before any others.

First: audit every publicly advertised job posting since 1 January 2026 and add the Ontario AI disclosure to any live or forthcoming posting immediately. The obligation under Ontario's regulation is in force now, applies to every posting, and attracts both regulatory and reputational risk if missed. Appoint one person — your HR lead or a named compliance contact — to own this check and sign off on every future posting before it goes live. This can be done within days and requires no external resource beyond the template in this pack.

Second: implement the Quebec Law 25 automated-decision notifications for your renewal-risk scoring and quote-pricing model without delay. Every straight-case quote issued without broker involvement must be accompanied by a clear statement that the decision was made by automated processing, identification of the principal factors used, and an explanation of the individual's right to request human review. This is the highest-penalty risk in your entire compliance picture — up to CAD $25 million or 4% of global turnover — and it has been a live legal obligation since September 2023. The Automated Decision-Making Policy in this pack gives you the operational framework; your priority is to ensure the notification is actually appearing in every relevant customer communication before the next quote cycle.

Third: convene a meeting of your senior leadership — ideally including your Chief Executive, Head of Operations and whoever is responsible for technology — to formally adopt the AI Register, assign an owner to each AI system in scope, and set a calendar of quarterly reviews. Governance on paper means nothing without named accountability. The OPC and CAI both look for evidence that a board or senior management team has actively engaged with AI risk, not merely received a document. Formalising ownership within the next 30 days demonstrates the good-faith governance posture that regulators treat as a mitigating factor in enforcement decisions.

Legal References

The following statutory and regulatory references underpin this summary. They are provided for your legal and compliance advisers and for any regulatory correspondence.

Personal Information Protection and Electronic Documents Act (PIPEDA), S.C. 2000, c. 5 — federal private-sector privacy statute governing all provinces except Quebec for federally regulated activities and all cross-border data flows.

Bill C-36 (tabled 15 June 2026) — proposed amendments to PIPEDA; not yet in force.

Quebec Act respecting the protection of personal information in the private sector (Law 25 / Bill 64) — automated decision-making notification and right to human review obligations; Privacy Impact Assessment requirement; maximum penalty CAD $25 million or 4% of global turnover.

Alberta Personal Information Protection Act (Alberta PIPA) — provincial privacy statute applicable to Beaumont Ridge's Alberta operations.

British Columbia Personal Information Protection Act (BC PIPA) — provincial privacy statute applicable to Beaumont Ridge's BC operations.

Ontario Employment Standards Act, 2000, and Ontario Regulation 476/24 — requirement for publicly advertised job postings to disclose use of AI in screening, assessment or selection of applicants; in force 1 January 2026; applies to employers of 25 or more.

Office of the Privacy Commissioner of Canada — guidance on artificial intelligence and privacy; active AI investigation programme as at August 2026.

Canada's National Artificial Intelligence Strategy: AI for All — published 4 June 2026; policy framework, not legislation.

Artificial Intelligence and Data Act (AIDA) — lapsed January 2025; not in force; not to be treated as a compliance obligation.

Provincial human rights codes — Ontario Human Rights Code; Quebec Charter of Human Rights and Freedoms; Alberta Human Rights Act; BC Human Rights Code — applicable to algorithmic discrimination claims arising from the renewal-risk scoring model.

PIPEDA Compliance Gap Analysis for AI Systems

Your AI use measured against the federal privacy principles.

Principle	Requirement	Current Position	Gap	Action	Priority
SUMMARY — Gap Count by Priority	Critical: 4 | High: 5 | Medium: 3 | Low: 1	As at 11 August 2026	—	See individual rows below	—
1. Accountability (PIPEDA Sched. 1, Principle 1)	A designated individual must be responsible for PIPEDA compliance, including all AI systems. Policies must name them.	No confirmed designated Privacy Officer on record for AI-specific governance across Toronto and Montreal offices.	No named accountable individual for AI processing; accountability chain for vendor AI tools (OpenAI, Microsoft) undocumented.	Appoint a named Privacy Officer covering all AI tools. Document accountability chain including vendor sub-processing agreements.	Critical
1. Accountability — Third-party vendors	PIPEDA s.4.1.3: organisations remain accountable for personal data transferred to processors.	Microsoft 365 Copilot and ChatGPT/OpenAI process policyholder and employee data. Vendor contracts not confirmed to address PIPEDA obligations.	No confirmed data processing agreements with Microsoft or OpenAI that bind them to PIPEDA-equivalent protections.	Execute or update DPAs with Microsoft and OpenAI. Confirm contractual prohibition on vendor model training on Beaumont Ridge client data.	Critical
2. Identifying Purposes (PIPEDA Sched. 1, Principle 2)	Purposes for collecting personal data must be identified before or at time of collection, including AI-specific uses.	Renewal-risk scoring model, claims triage tool, CRM, and HR AI each process personal data. Purposes not confirmed as formally documented per tool.	AI-specific processing purposes (scoring, triage, recruitment screening) not identified in a maintained purposes register.	Create a per-tool purposes register. Record that health and financial data are used in the renewal-risk model and claims triage tool.	High
3. Consent (PIPEDA Sched. 1, Principle 3) — General	Knowledge and consent required for collection, use, or disclosure of personal data, subject to exceptions.	Policyholders, claimants, and job applicants interact with AI tools. No evidence of AI-specific consent language in client-facing documentation.	Consent notices do not disclose AI processing of health and financial data for renewal scoring or claims triage.	Update policy documents, renewal notices, and client onboarding materials to disclose AI use and obtain meaningful consent.	Critical
3. Consent — Meaningful consent / non-obvious AI processing	OPC guidance (2023, 2025): consent is only meaningful if individuals understand that AI, not a human, is processing their data.	The renewal-risk model and claims triage tool operate without broker involvement on straight cases — not obvious to policyholders.	Individuals have no indication that an automated model sets their renewal price or routes their claim without human review.	Add plain-language disclosure at renewal and claim submission: 'An automated system processes this. You may request human review.'	Critical

Principle	Requirement	Current Position	Gap	Action	Priority
3. Consent — Quebec Law 25, Art. 12.1	Where a decision is made exclusively by automated processing, the individual must be informed and may request human review.	Renewal pricing on straight cases has no broker involved — this is an exclusively automated decision under Quebec Law 25 Art. 12.1.	No disclosed right to request human review for Quebec policyholders subject to automated renewal pricing.	Implement Quebec-specific notice at renewal. Provide a documented process for policyholders to request broker review of automated pricing.	Critical
3. Consent — Withdrawal	PIPEDA Sched. 1, Principle 3.8: individuals may withdraw consent subject to legal or contractual restrictions.	No documented process for policyholders or employees to withdraw consent to AI processing and receive a human alternative.	Withdrawal mechanism absent. Individuals cannot opt out of automated scoring or triage.	Establish a consent-withdrawal process. Where withdrawal would prevent service, document and disclose that restriction clearly.	High
4. Limiting Collection (PIPEDA Sched. 1, Principle 4)	Only collect personal data necessary for the identified purpose. AI models must not ingest superfluous fields.	Renewal-risk model and CRM AI ingest financial and health data on policyholders. Scope of fields fed to models not confirmed as minimised.	No confirmed data-minimisation review of inputs to the renewal-risk model, claims triage tool, or CRM AI.	Conduct a field-level audit of each AI tool's data inputs. Remove or pseudonymise fields not necessary for the stated purpose.	High
5. Limiting Use, Disclosure, Retention (PIPEDA Sched. 1, Principle 5)	Personal data must not be used or disclosed beyond the purpose for which it was collected, and must not be retained beyond that purpose.	No confirmed retention schedules specific to AI-processed data. Microsoft 365 Copilot may retain prompts and outputs containing client data.	Risk that AI-generated outputs (Copilot, ChatGPT) containing policyholder health or financial data are retained beyond the original purpose.	Set retention limits in Copilot and OpenAI tenancy settings. Include AI-generated outputs in the data retention and deletion schedule.	High
6. Accuracy (PIPEDA Sched. 1, Principle 6)	Personal data used to make decisions must be accurate, complete, and up to date.	The renewal-risk scoring model makes pricing decisions based on policyholder data held in the CRM and policy system.	No documented accuracy-assurance process to verify that model inputs are current and correct before a pricing decision is generated.	Implement a pre-scoring data-quality check. Provide policyholders a right to correct data and rerun the model if correction is material.	High
7. Safeguards (PIPEDA Sched. 1, Principle 7)	Personal data must be protected by security safeguards appropriate to its sensitivity. Health and financial data is high sensitivity.	Health and financial data processed by renewal model, claims triage, and CRM AI. Cloud vendors involved. Security controls not confirmed in gap analysis.	No confirmed AI-specific security assessment covering data in transit to and from Microsoft, OpenAI, and CRM vendor APIs.	Commission a security assessment covering API connections, prompt logging, and output storage for all six AI tools in scope.	High

Principle	Requirement	Current Position	Gap	Action	Priority
8. Openness (PIPEDA	Policies and practices regarding	Beaumont Ridge's public privacy policy	Privacy policy silent on AI processing,	Update the public privacy policy to list	Medium

Principle	Requirement	Current Position	Gap	Action	Priority
Sched. 1, Principle 8)	personal data management must be readily available to individuals.	does not appear to disclose the use of AI tools or the categories of decisions they affect.	automated renewal scoring, claims triage, and HR AI screening.	AI tools in use, the decisions they affect, and how individuals may seek more information.	
8. Openness — French-language obligation	Quebec Charter of the French Language: client-facing documents must be available in French. Beaumont Ridge serves Montreal clients.	Beaumont Ridge states it serves clients in French. AI disclosures and updated privacy policy must be available in French.	Any new AI consent notices or privacy policy updates will be deficient if not simultaneously published in French.	Translate all AI disclosure materials into French before publication. Apply this to renewal notices, claims notices, and the privacy policy.	Medium
9. Individual Access (PIPEDA Sched. 1, Principle 9)	Individuals have a right to access their personal data held by the organisation and to challenge its accuracy.	No confirmed process for policyholders to access data held in the renewal-risk model, CRM AI, or claims triage system.	Access request process does not appear to extend to AI-generated scores, model inputs, or triage decisions affecting an individual.	Extend the subject access process to cover AI-generated outputs and scores. Train staff to respond to access requests involving AI records.	Medium
10. Challenging Compliance (PIPEDA Sched. 1, Principle 10)	Individuals must be able to challenge compliance with PIPEDA. The designated individual handles complaints.	No AI-specific complaint route confirmed. The claims triage and renewal-pricing model make consequential decisions; complaints must be addressable.	Complaint process does not explicitly cover outcomes produced by automated tools or allow an individual to contest an AI decision.	Add an AI complaints pathway to the internal complaints procedure. Document how contested AI decisions are reviewed by a qualified broker.	Low
Ontario ESA / O. Reg. 476/24 — HR AI disclosure	In force 1 January 2026: public job postings must state if AI screens, assesses, or selects applicants. Applies to employers of 25+.	Beaumont Ridge (51–250 employees) posts vacancies publicly and uses an AI HR/recruitment system. Obligation is live.	Job postings as at 11 August 2026 not confirmed to include the required AI-screening disclosure.	Audit all active and future public job postings. Add compliant disclosure statement before next posting is published.	Critical
AI Bias — HR and Renewal Model	OPC guidance and PIPEDA Principle 6 (accuracy) require that AI decisions not be based on inaccurate or discriminatory proxies.	HR AI and renewal-risk model both make decisions affecting individuals. No bias audit confirmed for either tool.	Undetected proxy discrimination could affect renewal pricing (e.g., postal code as income proxy) or recruitment decisions.	Commission an independent bias audit of the renewal-risk model and HR AI. Document findings and corrective actions taken.	High
Federal AI legislation — AIDA / Bill C-36	AIDA (Bill C-27) lapsed January 2025 and will not be reintroduced. Bill C-36 was tabled 15 June 2026; it is not yet in force.	No standalone federal AI statute applies to Beaumont Ridge as at 11 August 2026. Obligations derive from PIPEDA and provincial law.	Bill C-36's privacy reforms are not yet enacted. Position should be monitored as C-36 progresses through Parliament.	Assign the Privacy Officer to track Bill C-36. Review the renewal-risk model and HR AI against C-36 draft provisions when available.	Medium

RATING SCALE — Critical: legal breach probable now, regulatory action or maximum fine exposure imminent, remediate within 30 days. High: material gap requiring remediation within 90 days. Medium: gap to address within 6 months, lower immediate enforcement risk. Low: process improvement, address within 12 months.

FINE EXPOSURE — Quebec Law 25 maximum penalty is CAD $25 million or 4% of worldwide turnover for the preceding fiscal year, whichever is greater. The turnover-based limb cannot be computed without Beaumont Ridge's financial figures. PIPEDA enforcement by the OPC results in public findings and Federal Court orders; monetary penalties under PIPEDA as currently in force are limited, but reputational and court-order risk is significant. Bill C-36, if enacted, would introduce higher penalties; position to be confirmed when the bill is enacted.

SCOPE OF AI TOOLS ASSESSED — (1) Automated renewal-risk scoring and quote-pricing model; (2) Claims triage routing tool; (3) AI-powered CRM / sales tools; (4) Microsoft 365 Copilot; (5) ChatGPT / OpenAI; (6) AI in HR / recruitment systems.

JURISDICTION NOTE — PIPEDA applies to Beaumont Ridge's commercial activities in Ontario, Alberta, and British Columbia. Quebec's Law 25 applies to Quebec operations and imposes the additional automated-decision notification obligation (Art. 12.1). BC PIPA and Alberta PIPA apply in those provinces and operate substantially in parallel with PIPEDA for private-sector organisations; where those statutes apply, the OPC's jurisdiction yields to the provincial commissioner, but the fair information principles assessed here are substantively equivalent across all four regimes. Ontario ESA / O. Reg. 476/24 is a separate employment-law obligation assessed in row 17 of this table for completeness, as it is triggered by Beaumont Ridge's AI-assisted recruitment.

COURT AND TRIBUNAL DUTIES — Where Beaumont Ridge files documents in any court or tribunal proceeding, counsel should confirm whether the relevant court rules or tribunal practice directions require disclosure that AI was used to draft or prepare the filing. This obligation is distinct from PIPEDA and is not governed by privacy law; it is noted here as a cross-cutting risk identified in the instructions.

REVIEW CADENCE — This gap analysis should be reviewed quarterly and updated whenever a new AI tool is onboarded, a material change is made to an existing tool's data inputs, or a new regulatory instrument (including Bill C-36) comes into force.

DOCUMENT OWNER — The designated Privacy Officer (to be appointed per row 1 action) is responsible for maintaining this gap analysis and reporting status to the Beaumont Ridge board on a quarterly basis.

Quebec Law 25 Compliance Roadmap

The Quebec-specific requirements, sequenced — including the ones with no equivalent elsewhere in the country.

1. Applicability of Quebec Law 25 to Beaumont Ridge Insurance Brokers

Quebec's Act respecting the protection of personal information in the private sector, as amended by Law 25 (Bill 64), applies in full to Beaumont Ridge Insurance Brokers. The firm operates a licensed office in Montreal, collects and processes personal information about Quebec-resident policyholders (including health and financial data), communicates in French, and uses automated systems — the renewal-risk scoring model and the claims-triage tool — that affect Quebec residents. All three phases of Law 25 are now in force. There is no threshold based on employee count; the law applies to any enterprise that collects, holds, uses or discloses personal information in carrying on an enterprise in Quebec. Beaumont Ridge meets every trigger.

☐ **Document in writing that Law 25 applies to Beaumont Ridge's Montreal operations and all Quebec-resident data subjects, and circulate that determination to senior leadership.**

The Chief Compliance Officer (CCO) drafts a one-page applicability memorandum citing ss. 3.1, 8 and 12 of the Act and the firm's Montreal office, Quebec licence, and Quebec-resident policyholder base. The memo is approved by the CEO and filed in the compliance register. It is reviewed annually or when a new Quebec-resident data processing activity begins.

Why it matters: Without a documented applicability determination, gaps in governance are invisible and the Commission d'accès à l'information (CAI) has no record that the firm has turned its mind to its obligations — which aggravates any penalty assessment up to CAD $25 million or 4% of global turnover.

2. Privacy Officer Designation and Publication

Section 3.1 of the Act (as amended by Law 25) requires every enterprise to designate a person responsible for the protection of personal information. That person's title and contact information must be published on Beaumont Ridge's website. The person responsible is, by default, the highest-ranking officer of the enterprise unless another person is designated. Given the firm's use of AI tools that process health and financial data, this role carries material accountability.

☐ **Formally designate a Privacy Officer by board or executive resolution and record the designation in writing.**

The CEO signs a designation letter naming the Privacy Officer (by title, e.g., Chief Compliance Officer or a dedicated Privacy Officer). The letter specifies the scope of responsibility — all personal information held by Beaumont Ridge in Quebec — and is retained in the corporate compliance register. If the role is held by an external consultant or shared-services provider, a written mandate must be in place.

Why it matters: Absence of a formal designation means the highest-ranking officer bears personal accountability by default, and the CAI may treat the gap as evidence of systemic non-compliance.

☐ **Publish the Privacy Officer's title and contact details on the Beaumont Ridge public website in both English and French.**

Add a dedicated 'Privacy / Confidentialité' page to the website displaying the Privacy Officer's title, a direct e-mail address or web form, and the firm's mailing address. The page must be reachable from the homepage footer. Because Beaumont Ridge serves clients in French, the page must appear in French as a minimum (s. 3.1 read with Quebec's Charter of the French Language); a bilingual page is best practice.

Why it matters: Failure to publish this information is a standalone violation the CAI can cite without any complaint from a data subject, and it undermines policyholder trust in a regulated sector where confidentiality is a selling point.

- [ ] **Owner and deadline: Privacy Officer / CCO to complete both steps by 30 September 2026.**

  The CCO coordinates with the Marketing / Web team for the website update and with the Corporate Secretary for the board resolution. Completion evidence (screenshot of live page, signed resolution) is filed in the compliance register.

  Why it matters: Both requirements have been in force since September 2023; any further delay extends existing exposure.

3. Automated Decision Transparency and the Right to Submit Observations

This is the requirement with no equivalent elsewhere in Canada. Section 12.1 of the Act requires that any person whose personal information is used to render a decision based exclusively on automated processing be informed of that fact and of the personal information used. That person has the right to be informed of and to submit observations on the decision, and to have a human review it. Beaumont Ridge's renewal-risk scoring and quote-pricing model operates 'with no broker involved on straight cases' — this is a textbook exclusively automated decision. The claims-triage routing tool may also qualify if routing determines access to or denial of a benefit without human intervention. Quebec Law 25 (s. 12.1) also requires notification when personal information is used for profiling.

- [ ] **Map every automated decision point in the renewal-risk / quote-pricing model and the claims-triage tool and classify each as 'exclusively automated' or 'human-in-the-loop'.**

  The CCO and IT lead conduct a process-walk of the renewal-risk model and claims-triage tool, documenting: (a) what personal information feeds each model; (b) what the output is (price, renewal/decline, claims routing); (c) whether a broker or adjuster reviews the output before it is communicated to the policyholder. Decisions where no human review occurs before communication are classified as exclusively automated and recorded in a Decision Inventory register.

  Why it matters: Misclassifying an exclusively automated decision as human-in-the-loop is the single most likely source of a CAI investigation for an insurer of this type.

- [ ] **Draft and implement a policyholder notification for each exclusively automated decision, delivered at the point the decision is communicated.**

  The notification must state in plain language: that the decision was made by automated processing; what categories of personal information were used (e.g., claims history, credit-based insurance score, health data); and that the individual has the right to request human review and to submit observations. For Quebec-resident policyholders the notice must be in French. Embed the notice in renewal letters, quote documents and claims acknowledgement messages. Reference s. 12.1 of the Act in the internal template notes.

  Why it matters: Omitting this notice exposes every automated renewal, quote and claims routing decision involving a Quebec resident as a separate breach of s. 12.1, each potentially attracting CAI investigation or administrative penalty.

- [ ] **Establish a documented human-review process that a Quebec-resident policyholder can invoke within a reasonable timeframe.**

  Define and document the review pathway: the policyholder submits a written request (by e-mail or letter) to the Privacy Officer; a qualified broker reviews the model output and the underlying data within 30 calendar days; the broker may confirm, vary or reverse the decision; and a written outcome is sent to the policyholder. The process is recorded in the policy manual and the Privacy Officer logs each request and outcome.

  Why it matters: Without an operational review pathway the right to submit observations is illusory, and the CAI is likely to treat it as a systemic violation rather than an isolated one.

- [ ] **Owner and deadline: CCO (decision inventory and process design) and VP Operations (broker review workflow) — complete by 31 October 2026.**

  CCO completes the Decision Inventory by 30 September 2026. VP Operations finalises the review workflow and trains affected brokers by 31 October 2026. The Privacy Officer publishes the review-request contact point on the website simultaneously with the Privacy Officer publication in Group 2.

  Why it matters: The obligation has been in force since September 2023; the firm's current 'no broker involved on straight cases' model is in active breach until remediated.

4. Privacy Impact Assessment (PIA) Duty

Section 3.3 of the Act requires a Privacy Impact Assessment (called a 'évaluation des facteurs relatifs à la vie privée' or EFVP in Quebec) before: (a) any acquisition, development or overhaul of an information system or electronic service delivery system involving personal information; (b) any communication of personal information outside Quebec; and (c) before entrusting personal information to a third party. For Beaumont Ridge, PIAs are triggered by: the renewal-risk / quote-pricing model (AI system involving health and financial data); the claims-triage tool; the AI-powered CRM; Microsoft 365 Copilot and ChatGPT / OpenAI (both of which may route data outside Quebec); the HR / recruitment AI system; and any future model changes. The CAI has published a guide on conducting EFVPs.

- [ ] **Conduct a retrospective PIA for each AI tool already in use that has not previously been assessed, prioritised by sensitivity of personal information processed.**

  Priority order: (1) renewal-risk / quote-pricing model (health and financial data, exclusively automated decisions); (2) claims-triage tool; (3) HR / recruitment AI (candidates' data); (4) AI-powered CRM; (5) Microsoft 365 Copilot and ChatGPT / OpenAI. For each, document: purpose and legal basis; data flows including cross-border transfers; risks to privacy; and mitigations implemented. Use the CAI's published EFVP guide as the framework. Retain completed PIAs for the life of the system plus five years.

  Why it matters: Operating an AI system involving personal information without a PIA is a violation of s. 3.3 and is an aggravating factor in any CAI investigation, particularly given that health data is involved.

- [ ] **Conduct a PIA before any new AI system or material change to an existing one is deployed.**

  Embed the PIA requirement in Beaumont Ridge's change-management and procurement processes: no AI tool or system involving personal information proceeds to deployment without a signed-off PIA. The Privacy Officer approves each PIA. Where ChatGPT / OpenAI or Microsoft 365 Copilot is used in a new use case, that use case must be assessed separately.

  Why it matters: A prospective PIA is the primary mechanism for demonstrating privacy-by-design, which the CAI treats as evidence of good faith in enforcement proceedings.

- [ ] **Complete a cross-border transfer assessment for Microsoft 365 Copilot and ChatGPT / OpenAI before continuing use involving Quebec-resident personal information.**

  Under s. 17 of the Act, personal information may only be communicated outside Quebec if a privacy assessment confirms that the recipient jurisdiction offers adequate protection. Document the assessment for each vendor, review their data processing agreements, and confirm or update contractual protections. If the assessment reveals inadequate protection, implement additional safeguards (e.g., pseudonymisation before upload, contractual clauses) or restrict Quebec-resident data from those tools.

  Why it matters: Sending Quebec-resident health or financial data to US-based AI services without a s. 17 assessment is one of the CAI's stated enforcement priorities and could expose the firm to penalties and mandatory remediation orders.

- [ ] **Owner and deadline: Privacy Officer owns PIA programme; retrospective PIAs for AI tools completed by 31 December 2026; cross-border assessments completed by 31 October 2026.**

  The Privacy Officer creates a PIA tracker listing each tool, assigned assessor, and target completion date. Progress is reported to the CEO quarterly. External counsel or a privacy consultant may be engaged for the cross-border assessments given their technical complexity.

  Why it matters: The obligation has applied since September 2023; a structured timeline with named owners converts existing exposure into a managed remediation programme that the CAI can verify.

5. Breach Reporting and the Confidentiality Incident Register

Sections 3.5 to 3.8 of the Act require Beaumont Ridge to: (a) maintain a register of all confidentiality incidents; (b) notify the CAI and affected individuals of any incident presenting a risk of serious injury; and (c) notify any third party to whom the information was communicated if that party can mitigate the harm. Given that Beaumont Ridge holds health and financial information on policyholders and routes data through ChatGPT / OpenAI and Microsoft 365 Copilot, breach risk is material. The notification threshold is a 'risk of serious injury' assessed by reference to the sensitivity of the information, the anticipated consequences, and the number of persons affected.

- ☐ **Create and maintain a confidentiality incident register covering all AI-processed personal information.**

  The register must record for each incident: the date discovered; a description of the incident; the personal information involved; the number of persons affected; whether notification was given and to whom; and corrective action taken. The register is held by the Privacy Officer, stored securely, and available for CAI inspection on request. Template entries are prepared in advance so that staff can populate them promptly.

  Why it matters: Failure to maintain the register is itself a violation of s. 3.5, independent of whether any reportable incident has occurred.

- ☐ **Establish an incident-response procedure with defined timelines for CAI notification and individual notification.**

  The procedure specifies: within 24 hours of discovery, the Privacy Officer is notified; within 72 hours, a preliminary risk assessment is completed; if a risk of serious injury is identified, the CAI is notified 'with the least possible delay' (no fixed statutory deadline, but the CAI expects prompt action); affected individuals are notified simultaneously or as soon as possible after the CAI. Notifications include the nature of the incident, the information involved, and steps taken. Brokers and IT staff receive annual incident-recognition training.

  Why it matters: Delayed notification after a breach involving policyholder health data would attract maximum regulatory scrutiny and reputational damage in the insurance market.

- ☐ **Owner and deadline: Privacy Officer drafts the register template and incident-response procedure by 30 September 2026; IT Security Lead integrates AI-tool monitoring alerts by 31 October 2026.**

  The Privacy Officer adapts existing PIPEDA breach protocols to meet the Quebec Law 25 standard (which is stricter). IT Security Lead configures monitoring for anomalous data access involving the renewal-risk model, CRM, and ChatGPT / OpenAI outputs. Both deliverables are tested with a tabletop exercise by 30 November 2026.

  Why it matters: An untested breach-response procedure is unlikely to meet the CAI's 'least possible delay' standard in a real incident.

6. Data Portability and De-indexing

Sections 28.1 and 28.2 of the Act grant Quebec residents the right to receive personal information they have provided to Beaumont Ridge in a structured, commonly used, technological format, and to have that information communicated to another enterprise where the information is processed by automated means (portability). Section 28.3 provides a right to de-indexing — the right to cease the dissemination of personal information or to re-index it — where the continued dissemination causes serious injury. These rights are exercisable against Beaumont Ridge in respect of policyholder data, CRM data, and data processed by AI tools. Portability applies to information that the individual has directly provided; de-indexing applies where information is disseminated on a technology-based medium.

- ☐ **Build and document a portability fulfilment process for Quebec-resident policyholders and HR candidates.**

  The Privacy Officer, working with IT, identifies the data sets that are: (a) provided by the individual; and (b) processed by automated means — these include CRM records, application data processed by the HR AI, and policyholder data fed into the renewal-risk model. A documented extraction process is created that outputs data in a structured, machine-readable format (e.g., CSV or JSON). Requests must be fulfilled within 30 days of receipt. A request form in French is published on the website alongside the Privacy Officer contact details.

  Why it matters: Inability to fulfil a portability request within a reasonable time is a direct violation of s. 28.1 and a likely catalyst for a CAI complaint in a competitive insurance market where switching is the point of the right.

- [ ] **Establish a de-indexing assessment protocol for any personal information about Quebec residents that Beaumont Ridge disseminates via a technology-based medium.**

  Audit all externally accessible platforms where individual personal information may appear (e.g., broker directories, testimonial pages, any AI-generated content). Create a documented process for receiving de-indexing requests: the Privacy Officer assesses whether the dissemination is causing serious injury; if yes, dissemination is ceased or re-indexing applied within 30 days; the outcome is recorded. If the request is refused, the individual is informed with reasons and their right to complain to the CAI.

  Why it matters: Given that Beaumont Ridge's AI tools may surface policyholder information in CRM outputs or internal knowledge-management responses, a failure to have a de-indexing process in place could expose health or financial information unexpectedly.

- [ ] **Owner and deadline: Privacy Officer and IT Lead to complete portability extraction process by 31 December 2026; de-indexing protocol by 30 November 2026.**

  IT Lead builds the extraction workflow; Privacy Officer drafts the request forms and response templates in both French and English. Both are tested with synthetic data before going live. Request forms are added to the website's Privacy page alongside the portability and de-indexing rights explanation.

  Why it matters: Both rights have been in force since September 2023; publishing rights without a functioning fulfilment process creates legal exposure if a request cannot be met.

7. Governance, Record-Keeping and Ongoing Obligations

Law 25 imposes several additional governance obligations that apply continuously to Beaumont Ridge: a published privacy policy in plain language (s. 3.2); consent requirements calibrated to the sensitivity of information (health and financial data require explicit consent); purpose-limitation (s. 5); and accuracy obligations (s. 11). The CAI is an active enforcement body; its maximum penalty of CAD $25 million or 4% of global turnover (whichever is greater) is the highest in Canada. The turnover-based limb cannot be computed here without Beaumont Ridge's confirmed global turnover figures.

- [ ] **Publish a plain-language, bilingual privacy policy on the Beaumont Ridge website that addresses AI-based processing, automated decision-making and cross-border transfers.**

  The Privacy Officer drafts the policy in both English and French. It must describe: categories of personal information collected; purposes of collection including AI-based uses; automated decision-making (s. 12.1 disclosure); cross-border transfers and the countries or regions involved; retention periods; and how to exercise rights (access, correction, portability, de-indexing, human review). The policy is reviewed annually and updated within 30 days of any material change to AI tools or data flows.

  Why it matters: A privacy policy that omits AI processing is inconsistent with s. 3.2 and undermines every other disclosure obligation in the Act.

- [ ] **Review and update consent mechanisms for health and financial data fed into the renewal-risk model and claims-triage tool.**

  The Privacy Officer and Legal counsel review the current consent language in policyholder application forms, renewal documents and claims forms. Where health or financial data is used in automated scoring, explicit and informed consent is required (s. 12 read with s. 12.1). Consent language must specifically identify automated processing and profiling as a purpose. Updated forms are deployed for all new Quebec-resident policies and renewals from the implementation date.

  Why it matters: Processing sensitive personal information without adequate consent is among the most serious violations under Law 25 and is directly in scope given that the renewal-risk model uses financial and potentially health data.

☐ **Assign a named owner and review date to every obligation in this roadmap and include the roadmap in the firm's annual compliance review cycle.**

The CCO creates a master compliance tracker cross-referencing each item in this roadmap to its statutory basis, named owner, target date and completion evidence. The tracker is reviewed at each quarterly compliance committee meeting. Completion evidence for CAI-facing obligations (PIA, privacy policy, automated decision notices) is retained for a minimum of five years.

Why it matters: Demonstrating a structured, documented compliance programme is the single most effective mitigant to the CAI's discretion to impose maximum penalties in the event of an incident or investigation.

CASL-Compliant AI Communications Procedures

Using AI to generate or send commercial messages without losing the consent position that makes them lawful.

1. Purpose and Scope

These procedures govern every commercial electronic message (CEM) sent or initiated by Beaumont Ridge Insurance Brokers in the course of its commercial and personal lines brokerage activities, where any part of that message — its content, personalisation, subject line, timing logic or recipient selection — is touched by an AI tool. The AI tools in scope include Microsoft 365 Copilot, ChatGPT / OpenAI, the AI-powered CRM and sales platform, and the automated renewal-risk scoring model where its outputs trigger or inform outbound communications. The claims triage tool and the HR / recruitment system are not normally used to generate CEMs directed at policyholders or prospects, but any future use of those tools to produce or initiate outbound messages to individuals will fall within this policy immediately and without amendment.

Canada's Anti-Spam Legislation (CASL, S.C. 2010, c. 23) applies to every CEM sent to an electronic address, including email, SMS and certain social-media messages, where the message has as one of its purposes encouraging participation in a commercial activity. Beaumont Ridge sends renewal notices, quote summaries, cross-sell offers, broker introductions and service communications — all of which are capable of constituting CEMs. The fact that a message is drafted or personalised by an AI tool does not alter its legal character; CASL applies to the message, not the drafter. These procedures are therefore mandatory for any staff member who uses an AI tool at any stage of a communications workflow, including prompt writing, template approval, list segmentation via CRM AI, and scheduled send functions.

2. Consent Types and How Each Is Evidenced

CASL section 6 prohibits sending a CEM unless the sender has either express or implied consent and the message contains the prescribed identification and unsubscribe information. Beaumont Ridge operates under both consent categories, and each must be evidenced in a discrete and retrievable record before an AI tool may be used to draft or send any message to that recipient.

Express consent exists where the individual has clearly and voluntarily agreed to receive CEMs from Beaumont Ridge, through a positive opt-in act such as ticking an unchecked box, signing a consent clause in a broker engagement letter, or completing an online quote request that includes a clearly worded consent statement. The consent record must capture the date and time of the opt-in act, the exact wording of the consent request presented to the individual, the channel through which consent was given, and the name of the staff member or system that recorded it. Express consent does not expire under CASL unless withdrawn, but Beaumont Ridge's CRM must flag any express consent record that has not been validated by a positive interaction in more than 36 months so that it can be reviewed before the next AI-generated message is sent.

Implied consent arises in defined circumstances under CASL sections 10(9) and 10(10), most relevantly where the individual has an existing business relationship with Beaumont Ridge — for example, a current or recently expired insurance policy placed through the firm — or where the individual has conspicuously published their electronic address and the message relates to their business role or functions. For existing policyholders, implied consent subsists for two years from the date the business relationship ends. For prospects who have made an inquiry in the preceding six months, implied consent subsists for six months from the date of inquiry. The CRM AI must be configured so that its segmentation and automated outreach logic reads the consent-expiry field before scheduling any message; the tool must not override or ignore that field. Where the CRM AI cannot confirm a valid consent basis for a contact, the contact must be excluded from the send list and the matter escalated to the Compliance Officer before any message is sent.

3. What AI-Generated Commercial Messages Must Still Contain

CASL section 6(2) sets out three categories of information that every CEM must contain regardless of how it was drafted. AI tools do not exempt Beaumont Ridge from any of them, and no AI-generated template may be approved for use unless all three are present in a non-removable format.

First, the message must clearly identify Beaumont Ridge Insurance Brokers as the sender. Where the CEM is sent on behalf of a specific broker or account manager, the message must identify both the individual and the firm. Microsoft 365 Copilot and ChatGPT drafts must be reviewed to confirm that the firm name has not been replaced by a generic sign-off or an AI-generated persona name. The CRM AI's personalisation tokens must be tested before each campaign to confirm they resolve correctly and do not produce blank or garbled sender fields.

Second, the message must contain contact information that enables the recipient to contact the sender, remaining valid for a minimum of 60 days after the message is sent. For Beaumont Ridge this means including a mailing address (Toronto or Montreal office address as appropriate), a telephone number or email address that is actively monitored, and, where the message is sent in connection with the Quebec book of business, ensuring the contact information is presented in French in compliance with the Charter of the French Language and the firm's Quebec licence obligations.

Third, the message must contain a functioning unsubscribe mechanism as described in section 4 below. AI drafting tools must not be permitted to suppress, reformat or relocate the unsubscribe block. Before any AI-generated message is sent at scale, a test send must confirm that the unsubscribe link resolves, that the landing page or reply mechanism functions, and that opt-out instructions are presented in both English and French where the recipient's preferred language is recorded as French or is unknown.

4. Identification and Unsubscribe Requirements That AI Generation Must Not Strip

The identification and unsubscribe elements are not stylistic — they are statutory minimums under CASL section 6(2)(a), (b) and (c), and their removal or degradation by an AI tool is a breach of the Act regardless of intent. These procedures impose the following hard controls.

Every CEM template used in conjunction with any AI tool must be stored in the CRM in a locked footer format that cannot be edited during the personalisation or drafting phase. The footer must contain: the full legal name of Beaumont Ridge Insurance Brokers, the registered office address for the province from which the communication originates, a monitored email address or telephone number, and the unsubscribe instruction. The unsubscribe instruction must use plain language — in English and in French where the recipient file records French as the preferred language — and must direct the recipient to a mechanism that processes the opt-out within ten business days as required by CASL section 11(3).

Where Copilot or ChatGPT is used to rewrite or shorten a message, the output must be pasted into the approved template; the AI output must never become the entire message including footer. Where the CRM AI generates and schedules messages automatically — for example, a renewal reminder triggered by the pricing model's identification of a straight-through renewal case — the system configuration must be audited quarterly to confirm that the locked footer is appended to every outbound message before delivery and that no AI personalisation step removes or truncates it. The Compliance Officer must certify each quarterly audit in writing. Any opt-out received through any channel — reply email, phone call, unsubscribe link, or verbal withdrawal recorded by the broker — must be entered into the CRM within two business days and must suppress all future AI-generated messages to that address immediately.

5. Record Keeping for Consent and for What Was Actually Sent

CASL places the burden of proof on the sender to demonstrate that consent existed at the time a CEM was sent (section 13). Where AI tools are involved in generating or dispatching messages, the record-keeping obligation extends to capturing not only the consent record but also the specific message that was actually delivered, since AI-generated content can vary between drafts and the approved version may differ from the sent version if human review steps are bypassed.

Beaumont Ridge must maintain the following records in a format that can be retrieved and produced to the Canadian Radio-television and Telecommunications Commission (CRTC) within a reasonable time on request: (a) the consent record for each recipient, including consent type, date, wording presented, and channel; (b) the consent expiry date and any renewal or refreshing events; (c) the name of the AI tool used to draft or personalise the message; (d) the final approved text of the message including all footer elements; (e) the date and time of send; (f) the from-address and the to-address; and (g) any bounce, opt-out or complaint record linked to that send.

Where the CRM AI or the renewal-pricing model triggers automated sends, the system must generate a send log that captures items (c) through (g) automatically. This log must not be overwritten or compressed within a minimum retention period of three years from the date of send. Where Copilot or ChatGPT is used interactively by a broker to draft a one-to-one CEM, the broker must copy the final sent text into the CRM contact record within one business day of sending. These records are also relevant to Beaumont Ridge's obligations under PIPEDA and Quebec Law 25, given that consent records contain personal information and the firm holds health and financial data on policyholders; access controls on these records must be consistent with those applied to policyholder data generally.

6. Review Before Send Where AI Drafted the Message

AI tools including Copilot and ChatGPT can produce text that is factually inaccurate, inconsistently personalised, or legally non-compliant. In a regulated financial services context, an inaccurate statement in a CEM about coverage, pricing or eligibility can expose Beaumont Ridge to liability under CASL, under provincial insurance legislation, and under the general law of misrepresentation. The firm's licences across Ontario, Quebec, Alberta and British Columbia require communications with clients and prospects to be accurate and not misleading. These review requirements are therefore both a CASL compliance control and a licensing obligation.

For bulk or campaign messages (defined as any send to more than five recipients generated or personalised using an AI tool), the following review steps are mandatory before the message is released: first, the drafting broker or marketing staff member must review the AI output against the source data — policy details, pricing outputs, CRM contact fields — and confirm accuracy; second, a second reviewer (either the Compliance Officer or a designated senior broker) must approve the final text and confirm that the statutory footer is present and functional; third, where the message is in French, a bilingual reviewer must confirm that the French text is accurate and that the unsubscribe instruction is correctly rendered in French. The two-reviewer sign-off must be recorded in the CRM task log linked to the campaign before the send is authorised.

For one-to-one AI-assisted messages, the drafting broker retains sole responsibility for reviewing the output before sending. Brokers must not use the send-on-behalf or scheduled-send functions within Copilot or the CRM without first completing a personal review of the full message as it will appear to the recipient. Where the renewal-risk scoring model or the claims triage tool produces an output that automatically triggers a client communication, that communication must pass through a human review queue and must not be sent without a broker's explicit approval, regardless of whether the case is classified as straight-through.

7. Liability Position for Directors and Officers

CASL section 31 provides that every officer, director, agent or mandatary of a corporation who directed, authorised, assented to, acquiesced in, or participated in the commission of a violation is personally liable for the violation, whether or not the corporation is proceeded against. CASL violations can attract administrative monetary penalties of up to CAD $1,000,000 for individuals and up to CAD $10,000,000 for organisations under section 20. CASL also provides a private right of action under section 47, which has been suspended by regulation but which may be activated; directors and officers should be aware that this provision exists and that its activation remains a legislative possibility.

For directors and officers of Beaumont Ridge, the practical implication is that delegating communications workflows to AI tools does not transfer legal responsibility away from the individuals who govern the firm. A director who approves a campaign send, who sets the parameters of the CRM AI's automated outreach logic, or who authorises the use of the renewal-pricing model's outputs to trigger unsolicited contact with prospects, is potentially exposed to personal liability if that send violates CASL. Directors and officers must therefore ensure that these procedures are formally adopted at board level, that compliance with them is reported to the board at least annually, and that any material departure from the review-before-send or consent-record requirements is escalated to the board immediately.

Officers responsible for technology and data (including any Chief Technology Officer, Chief Operating Officer, or equivalent) bear specific responsibility for the configuration of the AI tools and CRM systems described in these procedures. They must confirm in writing at least annually that the locked footer controls, consent-expiry logic, opt-out suppression, and send-log retention described in sections 3, 4 and 5 are functioning correctly. This written confirmation must be retained as part of the board's governance record. Directors and officers who discover a breach of these procedures must report it to the Compliance Officer within 24 hours; where the breach has resulted in CEMs being sent to individuals who had not consented or who had opted out, the Compliance Officer must assess whether self-reporting to the CRTC is appropriate and must seek legal advice within five business days of discovery.

AI Impact Assessment Template

Complete one before deploying any new AI system that touches personal data or affects people. Keep the completed copies.

This AI Impact Assessment Template is issued to Beaumont Ridge Insurance Brokers as part of its Canada AI Compliance — Complete pack. One completed copy of this form must be retained on file for every AI system that touches personal data or produces outputs that affect individuals before that system is deployed or materially changed. Completed assessments may be requested by the Office of the Privacy Commissioner of Canada (OPC), the Commission d'accès à l'information du Québec (CAI), or a counterpart provincial regulator. They are also relevant evidence in the event of a complaint, audit, or litigation. The template reflects obligations under PIPEDA, Quebec Law 25, BC PIPA, Alberta PIPA, Ontario ESA / O. Reg. 476/24, and OPC AI guidance current as at 11 August 2026. AIDA / Bill C-27 lapsed in January 2025 and will not be reintroduced; Bill C-36 was tabled on 15 June 2026 but is not yet in force. No obligation arising solely from C-36 is recorded here.

AI IMPACT ASSESSMENT
Beaumont Ridge Insurance Brokers
Version: [VERSION NUMBER] | Date completed: [DATE] | Next review: [REVIEW DATE] | Completed by: [FULL NAME AND TITLE] | Approved by: [SENIOR RESPONSIBLE OFFICER]

SECTION 1 — SYSTEM DESCRIPTION, PURPOSE AND PEOPLE AFFECTED

1.1 System name: [SYSTEM NAME — e.g. Renewal Risk Scoring and Quote-Pricing Model / Claims Triage Routing Tool / AI-Powered CRM / Microsoft 365 Copilot / ChatGPT-OpenAI Integration / HR Recruitment AI]

1.2 Vendor / provider: [VENDOR NAME AND REGISTERED ADDRESS]

1.3 Version or release assessed: [VERSION / RELEASE DATE]

1.4 Offices and jurisdictions where the system will operate: [SELECT ALL THAT APPLY — Toronto (Ontario) / Montreal (Quebec) / Alberta / British Columbia]

1.5 Purpose: Describe in plain language what the system does and the business outcome it produces.
[DESCRIBE PURPOSE — e.g. scores the likelihood that a commercial or personal lines policyholder will not renew, sets a quoted premium without broker involvement on straight cases, routes inbound claims to the correct handler, screens job applicants]

1.6 People affected: Identify every category of person whose data the system processes or whose interests an output affects.
[LIST CATEGORIES — e.g. policyholders (personal and commercial lines), claimants, prospective employees responding to public job postings, Beaumont Ridge staff whose performance data is analysed]

1.7 Estimated volume: [NUMBER OR RANGE OF INDIVIDUALS AFFECTED PER YEAR]

1.8 Does the system make or materially influence a decision about an individual without a broker or human reviewing that specific case before the decision is communicated? [YES / NO]

1.9 If yes, describe the decision and who receives it: [DESCRIBE — e.g. a renewal premium is issued to the policyholder; a claim is routed and a holding message sent; an applicant is screened out before any recruiter reviews the file]

SECTION 2 — NECESSITY AND PROPORTIONALITY

2.1 Why is an AI system necessary for this purpose rather than a manual or rule-based process?
[EXPLAIN — address speed, volume, consistency, and why human-only processing is insufficient or disproportionate]

2.2 Were less privacy-intrusive alternatives considered? State what they were and why they were rejected.
[LIST ALTERNATIVES CONSIDERED AND REASON FOR REJECTION]

2.3 Is the scope of personal data collected and processed the minimum necessary to achieve the stated purpose (data minimisation, PIPEDA Principle 4.4; Quebec Law 25 s. 12)?
[YES / NO — if no, describe what additional data is collected and justify it]

2.4 Does the system process special categories of information — specifically health or medical data, financial information, or information about minors — held by Beaumont Ridge?
[YES / NO — if yes, identify which categories and confirm they are required for the stated purpose]

Note: Beaumont Ridge holds health information on policyholders and financial information across its lines of business. Any system ingesting policy files or CRM records is likely to touch one or both categories. Confirm scope at 2.4 before proceeding.

2.5 Is the French-language service obligation engaged? Beaumont Ridge serves clients in French from its Montreal office. Confirm that system outputs, notices, and challenge routes are available in French to comply with Quebec language requirements.
[YES — CONFIRMED / NOT APPLICABLE TO THIS SYSTEM — EXPLAIN WHY]

SECTION 3 — DATA USED, LAWFUL BASIS AND RETENTION

3.1 List every category of personal data input to or generated by the system:
[INSERT TABLE OR LIST — e.g. name, address, policy number, claims history, health status indicators, credit or financial data, behavioural CRM data, CV and application data, performance metrics]

3.2 Source of each data category:
[e.g. policyholder application form, third-party data enrichment, internal CRM, HR system, Microsoft 365 activity logs]

3.3 Lawful basis for processing under PIPEDA and applicable provincial law:

PIPEDA: Processing personal information requires knowledge and consent (PIPEDA Schedule 1, Principle 3), unless a recognised exception applies. State the basis relied upon.
[CONSENT (form of consent obtained: [EXPRESS / IMPLIED]) / EXCEPTION RELIED UPON: [SPECIFY — e.g. PIPEDA s. 7(1)(b) where obtaining consent would compromise the availability or accuracy of the information and the collection is reasonable for purposes related to investigating a breach of an agreement]]

Quebec Law 25: Where the system is used in Quebec, confirm that the Privacy Policy disclosed to Quebec residents covers automated decision-making and profiling (Law 25 ss. 12, 63.1–63.3). Notification is required where a decision is made exclusively by automated processing.
[CONFIRMED / ACTION REQUIRED: [DESCRIBE]]

Alberta PIPA and BC PIPA: Consent or authority must be identified. State which applies.
[CONSENT / AUTHORITY — SPECIFY SECTION]

3.4 If the system produces outputs that exclusively determine an outcome affecting an individual (no human review), confirm that Beaumont Ridge has a mechanism to notify the individual that the decision was made exclusively by automated means, to explain the principal factors, and to offer a review by a person. This obligation arises under Quebec Law 25 and is consistent with OPC guidance.
[MECHANISM CONFIRMED: [DESCRIBE MECHANISM] / NOT YET IN PLACE — REMEDIATION PLAN: [DESCRIBE AND DATE]]

3.5 Retention period for personal data processed by or generated by the system:
[STATE RETENTION PERIOD AND LEGAL OR BUSINESS JUSTIFICATION — align with PIPEDA Principle 4.5 and CAI guidance]

3.6 Where is data stored, and does it leave Canada?
[DESCRIBE STORAGE LOCATION — e.g. Microsoft Azure Canada Central, OpenAI US servers]

3.7 If data is transferred outside Canada, what safeguards apply?
[e.g. contractual clauses, processor agreements, vendor sub-processor list reviewed — Quebec Law 25 requires a Privacy Impact Assessment before any transfer outside Quebec; confirm that assessment has been completed]

[SAFEGUARDS IN PLACE: [DESCRIBE] / CROSS-BORDER PIA COMPLETED: [YES / NO / DATE]]

SECTION 4 — RISKS TO INDIVIDUALS, SCORED, WITH MITIGATIONS

Score each risk: Likelihood (1 Low – 3 High) × Impact (1 Low – 3 High) = Raw Score (1–9). Record the residual score after mitigations.

4.1 Risk: Automated renewal-risk or pricing decision produces an outcome that discriminates on the basis of health status, postal code used as a proxy for ethnicity, or another protected ground under applicable human rights legislation.
Likelihood before mitigation: [1/2/3] | Impact: [1/2/3] | Raw score: [SCORE]
Mitigation: Conduct regular bias audits of model inputs and outputs; remove or justify any variable that correlates with a protected characteristic; document audit methodology and results; provide OPC and CAI with audit records on request.
Residual likelihood: [1/2/3] | Residual impact: [1/2/3] | Residual score: [SCORE]
Mitigation owner: [NAME / ROLE]

4.2 Risk: Claims triage routing tool misclassifies a claim, causing delay or denial of a legitimate claim, with no human review before the claimant is notified.
Likelihood before mitigation: [1/2/3] | Impact: [1/2/3] | Raw score: [SCORE]
Mitigation: Require human review of all triage outputs before external communication; set a confidence-threshold below which the tool escalates to a broker automatically; log all routing decisions for audit.
Residual likelihood: [1/2/3] | Residual impact: [1/2/3] | Residual score: [SCORE]
Mitigation owner: [NAME / ROLE]

4.3 Risk: Microsoft 365 Copilot or ChatGPT / OpenAI outputs incorporate client health or financial data into prompts that are retained by the vendor or used in model training, breaching client confidentiality and PIPEDA Principle 4.7 (safeguards).
Likelihood before mitigation: [1/2/3] | Impact: [1/2/3] | Raw score: [SCORE]
Mitigation: Enforce Microsoft 365 Copilot data boundary settings (commercial data protection enabled); implement an OpenAI API agreement with training-opt-out and data-deletion clauses; train staff not to paste personal or confidential data into consumer-facing AI interfaces; technical controls to prevent upload of restricted data categories.
Residual likelihood: [1/2/3] | Residual impact: [1/2/3] | Residual score: [SCORE]
Mitigation owner: [NAME / ROLE]

4.4 Risk: HR / recruitment AI screens out candidates without disclosure, breaching Ontario ESA O. Reg. 476/24 (in force 1 January 2026), which requires every publicly advertised job posting by an employer of 25 or more to state whether AI is used to screen, assess, or select applicants, including where a third party screens on Beaumont Ridge's behalf.
Likelihood before mitigation: [1/2/3] | Impact: [1/2/3] | Raw score: [SCORE]
Mitigation: Insert a standard disclosure statement in every public job posting: "Beaumont Ridge Insurance Brokers uses artificial intelligence tools to assist in screening, assessing, or selecting applicants for this role"; audit all third-party recruitment platforms to confirm the same disclosure appears where they post on Beaumont Ridge's behalf; retain posting records as evidence of compliance.
Residual likelihood: [1/2/3] | Residual impact: [1/2/3] | Residual score: [SCORE]
Mitigation owner: [NAME / ROLE]

4.5 Risk: Performance management data processed through AI-powered CRM or HR tools influences employment decisions (discipline, termination, promotion) without the employee being aware that AI contributed to the assessment.
Likelihood before mitigation: [1/2/3] | Impact: [1/2/3] | Raw score: [SCORE]
Mitigation: Disclose in the employee privacy notice that AI tools may process performance data; ensure no employment decision is communicated as AI-only; provide a human review route before any adverse outcome.
Residual likelihood: [1/2/3] | Residual impact: [1/2/3] | Residual score: [SCORE]
Mitigation owner: [NAME / ROLE]

4.6 Risk: Breach or unauthorised access to personal data held within the AI system triggers mandatory breach notification obligations under PIPEDA s. 10.1 (real risk of significant harm threshold) and Quebec Law 25 s. 63.3 (notification to CAI and affected individuals).
Likelihood before mitigation: [1/2/3] | Impact: [1/2/3] | Raw score: [SCORE]
Mitigation: Encrypt data at rest and in transit; conduct penetration testing before deployment; maintain a breach response plan naming the Privacy Officer; confirm vendor incident notification SLA is 24 hours or less.
Residual likelihood: [1/2/3] | Residual impact: [1/2/3] | Residual score: [SCORE]
Mitigation owner: [NAME / ROLE]

4.7 Risk: [ANY ADDITIONAL RISK IDENTIFIED DURING ASSESSMENT]
Likelihood before mitigation: [1/2/3] | Impact: [1/2/3] | Raw score: [SCORE]
Mitigation: [DESCRIBE]
Residual likelihood: [1/2/3] | Residual impact: [1/2/3] | Residual score: [SCORE]
Mitigation owner: [NAME / ROLE]

SECTION 5 — HUMAN OVERSIGHT AND ROUTE TO CHALLENGE

5.1 Human oversight model: Describe the point at which a qualified person reviews, confirms, or can override the system's output before it has effect.
[DESCRIBE — e.g. broker reviews flagged renewals above a premium-change threshold before issuance; claims handler reviews triage output before sending holding acknowledgement to claimant]

5.2 Fully automated decisions: If any decision is communicated to the individual without prior human review, state which decisions these are and confirm that the automated-decision disclosure mechanism described at 3.4 is in place.
[LIST AUTOMATED DECISIONS / CONFIRM MECHANISM OR STATE REMEDIATION DATE]

5.3 Individual's right to explanation: Describe how an affected individual (policyholder, claimant, job applicant, employee) can request an explanation of an AI-influenced outcome. This is required under Quebec Law 25 where a decision is made exclusively by automated processing, and is consistent with OPC guidance for high-impact systems across all jurisdictions.
[DESCRIBE PROCESS — include contact point, language (English and French), and response timeframe]

5.4 Route to challenge or request human review: Describe how an individual can challenge an outcome and request that a person, rather than the system, determines the matter.
[DESCRIBE — e.g. the individual contacts the Complaints Officer by email at [EMAIL ADDRESS]; the Complaints Officer must acknowledge within [X] business days and provide a human-reviewed determination within [X] business days; the outcome is recorded in the CRM log]

5.5 Escalation pathway within Beaumont Ridge:
[MAP — e.g. frontline broker -> Compliance Officer -> Privacy Officer -> Chief Operating Officer -> Board]

5.6 Court and tribunal disclosure: Where Beaumont Ridge files documents in legal proceedings, confirm there is a policy requiring disclosure of any AI-generated or AI-assisted content in accordance with applicable court or tribunal directions.
[POLICY IN PLACE: YES / NO — if no, record action and owner: [NAME] by [DATE]]

SECTION 6 — RESIDUAL RISK, SIGN-OFF AND REGULATOR CONSULTATION

6.1 Summary of residual risk after all mitigations:
[SUMMARISE — list any risk that remains at residual score 6 or above and explain why it is accepted or what further action is planned]

6.2 Overall residual risk rating: [LOW / MEDIUM / HIGH]

6.3 Deployment decision:
[ ] Approved for deployment — all mitigations confirmed in place
[ ] Conditionally approved — deployment permitted from [DATE] subject to mitigations at items [LIST] being completed by [DATE]

[ ] Not approved — system must not be deployed until residual risk is reduced

6.4 Sign-off:
Privacy Officer: [NAME] | Signature: _____________ | Date: _____________
Compliance Officer: [NAME] | Signature: _____________ | Date: _____________
Senior Responsible Officer: [NAME] | Signature: _____________ | Date: _____________

6.5 Must the regulator be consulted before deployment?

OPC: OPC guidance recommends that organisations conducting high-impact AI processing proactively engage the OPC, particularly where health or financial data is processed at scale. Formal pre-deployment consultation is not a statutory requirement under PIPEDA as currently in force, but is strongly advisable where this assessment identifies residual risks scored 6 or above. PIPEDA s. 10 requires that Privacy Management Programme records be available to the OPC on request.
[CONSULTATION REQUIRED: YES / NO — if yes, record date of approach and OPC reference: [REFERENCE]]

CAI (Quebec): Where the system is deployed in Quebec and involves automated decision-making or transfer of personal information outside Quebec, a Privacy Impact Assessment must be submitted to the CAI before deployment (Quebec Law 25 s. 63.5 and related provisions). Maximum penalty for breach: CAD $25,000,000 or 4% of worldwide turnover (the higher applies; the turnover-based limb cannot be computed without Beaumont Ridge's figures).
[CAI PIA COMPLETED: YES / DATE / REFERENCE NUMBER — or — NOT YET COMPLETED: TARGET DATE [DATE], OWNER [NAME]]

Alberta OIPC and BC OIPC: Neither Alberta PIPA nor BC PIPA imposes a mandatory pre-deployment consultation, but both commissioners may request records. Confirm that this assessment is filed and retrievable.
[FILED: YES / NO]

SECTION 7 — REVIEW DATE AND VERSION CONTROL

7.1 This assessment must be reviewed and, where necessary, updated:
— before any material change to the system's purpose, data inputs, or decision logic;
— following any data breach or regulatory inquiry involving the system;
— at the scheduled review date below, and in any event not less than annually.

7.2 Scheduled review date: [DATE — no more than 12 months from the date at 6.4]

7.3 Version history:

Version [NUMBER] | Date: [DATE] | Author: [NAME] | Change summary: [DESCRIPTION — e.g. Initial issue / Updated following bias audit / Updated to reflect vendor change]

[ADD ROWS AS REQUIRED]

7.4 Document owner: [ROLE — e.g. Privacy Officer, Beaumont Ridge Insurance Brokers]

7.5 Storage location of completed assessment: [LOCATION — e.g. SharePoint compliance library, restricted access, retained for minimum [X] years in accordance with Beaumont Ridge document retention policy]

END OF TEMPLATE

Provincial Applicability Assessment

Which provincial privacy regimes reach this business, and what each changes.

Province	Regime	Applies	Basis	What it changes for Beaumont Ridge
Federal (all provinces)	PIPEDA (as current, pre-Bill C-36 enactment)	Yes	Beaumont Ridge collects, uses and discloses personal information in the course of commercial activities across all four licensed provinces.	Requires consent, purpose limitation, safeguards and breach notification. OPC AI guidance applies to renewal-scoring model, claims triage tool and CRM. Bill C-36 tabled 15 June 2026 — not yet in force; position to be confirmed on enactment.
Quebec	Quebec Law 25 (Act 25, amending Act respecting the protection of personal information in the private sector)	Yes	Beaumont Ridge operates a Montreal office, holds Quebec policyholder data including health and financial information, and serves clients in French.	Automated-decision transparency: where renewal-risk scoring or claims triage makes a decision exclusively by automated processing, Beaumont Ridge must notify the individual and, on request, explain the decision and allow human review. Profiling rules apply to CRM segmentation. Maximum fine: CAD $25 m or 4% global turnover.
Quebec	Quebec Law 25 — language obligation	Yes	Beaumont Ridge serves Quebec clients in French as declared.	Privacy notices, consent forms, automated-decision explanations and any AI-use disclosures directed at Quebec policyholders must be available in French. English versions alone are non-compliant.
Ontario	PIPEDA (federal) — substantially similar status not granted; Ontario has no private-sector privacy statute of general application	Federal PIPEDA governs	Ontario has not enacted substantially similar private-sector privacy legislation. PIPEDA applies directly to Beaumont Ridge's Ontario commercial activities.	No Ontario-specific private-sector privacy additions beyond PIPEDA. Ontario ESA / O. Reg. 476/24 applies separately to HR use — see HR row.
Ontario	Ontario ESA / O. Reg. 476/24 — AI in recruitment	Yes	Beaumont Ridge has 51–250 employees (25+ threshold met) and posts vacancies publicly. In force 1 January 2026.	Every publicly posted job must state whether AI is used to screen, assess or select applicants, including where a third-party HR/recruitment AI system performs that function on Beaumont Ridge's behalf. Non-disclosure is a breach of the ESA.
Alberta	Alberta PIPA (Personal Information Protection Act)	Yes	Beaumont Ridge is licensed in Alberta and collects personal information from Alberta policyholders in the course of commercial activities.	Substantially similar to PIPEDA; PIPEDA does not apply to provincially regulated activities in Alberta. Consent, access and correction rights governed by Alberta PIPA. OPC oversight replaced by Alberta OIPC. Automated-decision rules are less prescriptive than Quebec Law 25 — PIPEDA/Alberta PIPA floor applies; Quebec standard adopted as default (see notes).

Province	Regime	Applies	Basis	What it changes for Beaumont Ridge
British Columbia	BC PIPA (Personal Information Protection Act)	Yes	Beaumont Ridge is licensed in BC and collects personal information from BC policyholders in the course of commercial activities.	Substantially similar to PIPEDA; PIPEDA does not apply to provincially regulated activities in BC. Consent, access and correction rights governed by BC PIPA. OPC oversight replaced by BC OIPC. No express automated-decision notification obligation equivalent to Quebec Law 25 — Quebec standard adopted as default.
Alberta & BC	Employee personal information — Alberta PIPA / BC PIPA	Yes	Both statutes govern employee personal information collected by private-sector employers in those provinces, including for performance management and HR AI tools.	Beaumont Ridge's AI-powered HR/recruitment and performance-management systems must comply with Alberta and BC PIPA employee provisions: purpose limitation, notice, and safeguards. No separate health information statute displaces this for employee records in either province.
Quebec	Employee personal information — Quebec Law 25	Yes	Quebec Law 25 covers employee personal information held by private-sector organisations, including at Beaumont Ridge's Montreal office.	HR/recruitment AI and performance-management tools processing Montreal employee data must comply with Law 25 consent, automated-decision transparency and retention rules. Profiling of employees via AI CRM or HR tools triggers notice obligations.
Ontario	Employee personal information — Ontario has no private-sector employee privacy statute; PIPEDA governs	Federal PIPEDA governs	Ontario did not enact substantially similar legislation covering employee personal information in the private sector.	PIPEDA applies to Beaumont Ridge's Ontario employee data. No additional provincial layer beyond the ESA AI-disclosure obligation for recruitment postings.
Quebec	Health information — Quebec Law 25 / Act respecting health and social services information	Yes — Law 25 applies; sector-specific health statute governs public-sector health data, not private insurers	Beaumont Ridge holds policyholder health information as a private-sector insurer/broker. Quebec Law 25 governs as the applicable private-sector regime.	Policyholder health data is sensitive personal information under Law 25. Heightened consent, strict purpose limitation, and breach notification obligations apply. Automated use of health data in renewal-scoring model requires explicit justification and automated-decision disclosure.
Alberta	Health information — Alberta Health Information Act (HIA)	No — HIA governs custodians (clinics, hospitals, etc.), not insurance brokers	Beaumont Ridge is not a health information custodian under the HIA. Policyholder health data held for underwriting purposes falls under Alberta PIPA.	Alberta PIPA governs health-related personal information collected by Beaumont Ridge. No separate HIA obligations arise. Sensitive-data safeguard expectations under PIPA still apply.
BC	Health information — BC does not have a	No separate regime — BC PIPA governs	BC's health information rules (e.g., E-Health Act) target public-sector and	Policyholder health data held by Beaumont Ridge in BC is governed by BC PIPA. Sensitive-data safeguards

Province	Regime	Applies	Basis	What it changes for Beaumont Ridge
	private-sector health information act separate from BC PIPA for insurers		health-authority data, not private insurance brokers.	and purpose-limitation obligations apply. No additional layer.
Ontario	Health information — Ontario PHIPA	No — PHIPA governs health information custodians; Beaumont Ridge is not one	Beaumont Ridge is an insurance broker, not a health information custodian under PHIPA (hospitals, practitioners, etc.).	Policyholder health data held for underwriting/claims falls under PIPEDA, not PHIPA. No PHIPA obligations arise. Sensitive-data protections under PIPEDA apply.
Strictest applicable position — default standard	Quebec Law 25	Adopted as default across all operations	Quebec Law 25 imposes the most onerous automated-decision, consent, and breach-notification obligations, backed by the highest penalty: CAD $25 m or 4% global turnover.	Beaumont Ridge should apply Quebec Law 25 standards enterprise-wide: automated-decision notice, human-review rights, French-language compliance, and sensitive-data safeguards. This satisfies PIPEDA, Alberta PIPA and BC PIPA simultaneously.

SUBSTANTIALLY SIMILAR LEGISLATION. Alberta PIPA and BC PIPA are recognised by the federal government as substantially similar to PIPEDA for provincially regulated private-sector activities. PIPEDA is therefore displaced by those statutes for Beaumont Ridge's activities in Alberta and BC that fall within provincial jurisdiction. Quebec Law 25 is likewise substantially similar. Ontario has no substantially similar private-sector privacy statute; PIPEDA applies directly and in full to Ontario operations. | AUTOMATED DECISION-MAKING. Quebec Law 25 is the only applicable regime with an express statutory obligation to notify individuals where a decision is made exclusively by automated processing and to provide a human-review mechanism on request. Beaumont Ridge's renewal-risk scoring and quote-pricing model (which operates with no broker involvement on straight cases) and the claims triage routing tool both fall within scope. These tools must be documented, and Quebec policyholders must receive notification and, on request, an explanation and the right to have a human reconsider the outcome. | FEDERAL AI LEGISLATION. The Artificial Intelligence and Data Act (AIDA), included in Bill C-27, lapsed in January 2025 and has been confirmed as discontinued. Bill C-36 (tabled 15 June 2026) revives private-sector privacy reforms but is not yet in force. No standalone federal AI statute is currently in force. Canada's National AI Strategy (AI for All, launched 4 June 2026) is a policy framework, not binding legislation. Beaumont Ridge should monitor Bill C-36 progress; the position should be confirmed on enactment. | OPC GUIDANCE AND INVESTIGATIONS. The OPC is actively issuing AI-related privacy guidance and conducting investigations. OPC guidance does not carry the force of statute but reflects how the Commissioner will assess PIPEDA compliance. Beaumont Ridge's use of ChatGPT/OpenAI and Microsoft 365 Copilot for financial analysis and internal knowledge management should be reviewed against OPC AI and privacy guidance to confirm personal data is not transferred to third-party AI providers without adequate contractual safeguards. | RECRUITMENT DISCLOSURE — ONTARIO. O. Reg. 476/24 under the Employment Standards Act came into force 1 January 2026. Every public job posting by Beaumont Ridge (51–250 employees, well above the 25-employee threshold) must state whether AI is used to screen, assess or select applicants. This obligation extends to any third-party recruitment AI platform acting on Beaumont Ridge's behalf. Existing postings that pre-date 1 January 2026 should be reviewed; all new postings must include the disclosure. | FRENCH-LANGUAGE OBLIGATION. All privacy notices, consent mechanisms, automated-decision explanations, and AI-use disclosures directed at Quebec individuals must be available in French. This applies to policyholder-facing chatbot interactions and CRM communications, as well as any written explanation provided following a request under Law 25's automated-decision provisions. | PENALTY SCALE. Quebec Law 25: up to CAD $25 million or 4% of worldwide turnover for the most serious violations. The turnover-based limb cannot be computed here without Beaumont Ridge's audited global revenue figures. Alberta PIPA and BC PIPA: maximum fines are lower (up to CAD $100,000 per organisation under each statute). PIPEDA: OPC can refer matters to the Federal Court; fines under PIPEDA's breach-notification provisions up to CAD $100,000. | REVIEW CADENCE. This assessment should be reviewed: (a) on enactment or material amendment of Bill C-36; (b) on any OPC enforcement decision or guidance affecting insurance sector AI use; (c) annually as a minimum given the pace of regulatory change. Owner: Chief Compliance Officer, with Legal counsel input for Quebec Law 25 automated-decision obligations.

Federal AI Direction and Readiness Review

Where federal AI regulation actually stands, and what to build now that is required regardless of what Parliament does next.

1. The Current Status of Federal AI Legislation in Canada

Beaumont Ridge Insurance Brokers operates at a moment when federal AI-specific legislation does not exist in Canada and is not imminent. Bill C-27, which carried the Artificial Intelligence and Data Act (AIDA), lapsed in January 2025 and the federal government has confirmed that AIDA will not be reintroduced. There is therefore no federal AI statute binding Beaumont Ridge today, and no single Canadian federal law that imposes AI-specific obligations on private-sector organisations.

On 15 June 2026, the federal government tabled Bill C-36. Bill C-36 revives the private-sector privacy reforms that had been embedded within C-27 and would amend PIPEDA, but it does not reintroduce AI-specific legislation in the manner AIDA contemplated. Bill C-36 is at introduction stage and has not received Royal Assent; its provisions are not yet law. The position of Bill C-36 must be monitored, as amendments to PIPEDA could affect the lawful-basis and transparency framework that governs Beaumont Ridge's AI-driven processing of personal information.

Canada's National Artificial Intelligence Strategy, titled AI for All, was launched on 4 June 2026. It is a five-year policy framework, not a statute. It sets a target of raising AI adoption among Canadian businesses from 12 per cent to 60 per cent by 2034 and commits C$925.6 million from Budget 2025 to AI infrastructure. The Strategy signals the government's intent to govern AI primarily through privacy law reform, online safety measures and investment rather than through standalone AI legislation. Beaumont Ridge should treat the Strategy as an indicator of regulatory direction, not a source of current legal obligation.

In the absence of AIDA or any successor, the Office of the Privacy Commissioner of Canada (OPC) is the principal federal body shaping AI governance expectations. The OPC has issued guidance on AI and privacy and, in May 2026, published draft age assurance guidance. OPC guidance is not legally binding in the same sense as a statute, but the OPC can investigate complaints, issue findings and impose orders under PIPEDA; its guidance documents define the standard of conduct against which Beaumont Ridge's practices will be measured in any investigation. Given that the OPC has active AI investigations under way, Beaumont Ridge should treat current OPC guidance as effectively mandatory.

2. Binding Obligations Today: PIPEDA and Provincial Privacy Regimes

PIPEDA applies to Beaumont Ridge's commercial activities involving personal information across all four provinces in which it is licensed. PIPEDA's ten fair information principles, set out in Schedule 1, require that personal information is collected only for identified purposes (Principle 2), with knowledge and consent (Principle 3), and that individuals are informed of those purposes at or before collection (Principle 8). Where Beaumont Ridge's automated renewal-risk scoring model and claims triage tool process policyholder data — including health and financial information — to make or influence decisions, those processes must be disclosed to individuals as a purpose of collection.

Quebec Law 25 imposes additional obligations that apply to Beaumont Ridge's Montreal office and to any Quebec residents whose data is processed. Under section 12.1 of the Act respecting the protection of personal information in the private sector (as amended by Law 25), an enterprise must inform an individual when a decision is made based exclusively on automated processing of personal information, must inform the individual of the personal information used, must allow the individual to present observations to a person able to review the decision, and must, on request, allow the individual to be re-examined by a human. This requirement is directly triggered by Beaumont Ridge's renewal-risk scoring and quote-pricing model when it operates without broker involvement on straight cases. Quebec's maximum penalty is CAD $25 million or four per cent of worldwide turnover, whichever is greater; the turnover-based limb cannot be computed here without Beaumont Ridge's financial figures, but the absolute maximum is CAD $25 million.

Alberta PIPA and British Columbia PIPA apply to Beaumont Ridge's activities in those provinces. Both statutes follow a consent-and-purpose framework comparable to PIPEDA and require that individuals be told of the purposes for which personal information is collected before or at the time of collection. Neither statute currently contains an explicit automated-decision-making notification provision equivalent to Quebec's section 12.1, but both require meaningful consent and transparency. Where health or financial information — both of which are sensitive under each provincial statute — is fed into Beaumont Ridge's AI models, a higher standard of express consent applies. Beaumont Ridge should confirm the current position under Alberta PIPA and BC PIPA with qualified provincial counsel, as the briefing indicates further amendments to PIPEDA in 2026 that could affect interpretation.

Ontario does not have a private-sector privacy statute separate from PIPEDA, so PIPEDA governs personal information processed in connection with Beaumont Ridge's Ontario operations. Ontario's Employment Standards Act and O. Reg. 476/24 apply specifically to AI in recruitment and are addressed in section 5 below.

3. High-Impact System Criteria and Assessment of Beaumont Ridge's AI Tools

Although AIDA did not come into force, the concept of a high-impact AI system — defined in the AIDA draft regulations as a system that, among other characteristics, makes or assists in making consequential decisions about individuals in domains including financial services, employment and insurance — remains the organising principle adopted by the OPC and reflected in Quebec's Law 25 approach. Beaumont Ridge should apply that criteria framework now because it reflects what the OPC will apply in an investigation and what any successor legislation is likely to codify.

The automated renewal-risk scoring and quote-pricing model meets every criterion for a high-impact system. It processes health and financial information about policyholders; it produces outputs that directly determine whether a policyholder's coverage is renewed and at what premium; and on straight cases it operates without a broker making an independent assessment. This is a consequential automated decision about individuals in the insurance domain. The absence of human review on those cases brings the model squarely within the Quebec Law 25 section 12.1 notification and review entitlement.

The claims triage routing tool processes personal information to route claims, which determines the speed and pathway of a claim's resolution. Routing decisions can affect whether a claimant receives urgent attention or is delayed. The tool should be treated as high-impact because an incorrect routing decision carries material financial and practical consequences for the claimant.

The AI in HR and recruitment systems processes personal information about job applicants and employees to assist in screening, assessment, performance management or related decisions. Decisions about hiring, promotion and performance management are consequential decisions about individuals and engage both privacy obligations and, in Ontario, specific disclosure duties under O. Reg. 476/24.

Microsoft 365 Copilot and ChatGPT / OpenAI are general-purpose productivity tools. Their impact level depends on use. Where they are used for internal drafting or knowledge management without processing identifiable personal information, they present lower risk. Where they are used to process policyholder data — for example, summarising claim files or drafting communications referencing identifiable clients — they must be treated as processing personal information under PIPEDA and provincial statutes, and data-sharing terms with the vendors must be reviewed to confirm sub-processing arrangements meet consent and purpose-limitation obligations.

The AI-powered CRM and sales tools process personal information about prospects and clients in the course of commercial activities. Their risk level is moderate to high: they may profile individuals, score sales opportunities and influence which clients receive which products or services. Profiling of individuals in a commercial context requires transparent disclosure under PIPEDA and Quebec Law 25.

4. Accountability Framework, Mitigation and Monitoring Duties

PIPEDA Principle 1 requires that an identifiable individual at Beaumont Ridge be accountable for the organisation's compliance with the fair information principles, and that policies and practices be in place to give effect to them. Beaumont Ridge must designate a Chief Privacy Officer or equivalent role with explicit responsibility for AI-related processing. That individual's responsibilities must be documented, and contact information must be made available to individuals on request.

For each high-impact AI system identified above, Beaumont Ridge must conduct and document a Privacy Impact Assessment (PIA). The OPC has confirmed through its AI guidance that PIAs are expected before deployment of AI systems that process personal information to make decisions affecting individuals. Each PIA must identify the personal information processed, the lawful basis for processing, the risks of inaccurate or biased outputs, the measures taken to mitigate those risks, and the retention and deletion arrangements. For the renewal-risk scoring model, the PIA must specifically address whether the model produces differential outcomes across protected groups — for example, whether health data inputs correlate with characteristics protected under provincial human rights legislation — and document the steps taken to test and remediate bias.

Quebec Law 25 requires a Privacy Impact Assessment before any new information system involving personal information is put into service or before communicating personal information outside Quebec. Beaumont Ridge's use of ChatGPT / OpenAI and potentially Microsoft 365 Copilot involves transferring personal information to servers outside Canada. A Law 25 PIA for those data flows is mandatory, and Beaumont Ridge must satisfy itself that the level of protection is equivalent to that required under Quebec law before enabling those tools to process Quebec residents' personal information.

Bias and discrimination risk is a live concern across three of Beaumont Ridge's AI systems. The renewal-risk scoring model uses health and financial information; in the insurance context, pricing and renewal decisions that correlate with health status can engage provincial human rights codes in Ontario, Quebec, Alberta and British Columbia. Beaumont Ridge must implement pre-deployment bias testing using representative datasets, document the results, and repeat testing at defined intervals — at minimum annually and whenever the model is retrained or its inputs change materially. A human review process must be established for cases where the model's output is challenged by a policyholder.

The OPC's active AI investigations are a material risk factor. Beaumont Ridge should treat any OPC investigation request as a trigger for full documentation disclosure; incomplete records are treated as evidence of non-compliance. Monitoring must therefore be ongoing: Beaumont Ridge must log model decisions, retain audit trails sufficient to reconstruct how a specific output was reached for a specific individual, and maintain records of bias-testing results. Monitoring cadence must be formally scheduled, not ad hoc.

5. Record-Keeping and Publication Duties

PIPEDA Schedule 1, Principle 1 requires that policies and practices implementing the fair information principles are documented and that the organisation makes information about those policies and practices available on request. Beaumont Ridge must maintain a documented AI use register that identifies each AI tool in use, its purpose, the personal information it processes, the decisions it supports or makes, and the accountability owner for each.

Under Quebec Law 25, Beaumont Ridge must publish a privacy policy that is clear, plain-language and available on its website. Because Beaumont Ridge serves clients in French and operates in Quebec, that policy must be available in French. The policy must describe how personal information is used, including its use in automated decision-making, and how individuals can exercise their rights including the right to request human review of automated decisions under section 12.1.

Ontario O. Reg. 476/24, in force from 1 January 2026, requires that any publicly advertised job posting disclose whether artificial intelligence is used to screen, assess or select applicants. Beaumont Ridge has 51 to 250 employees — well above the 25-employee threshold — and posts its own vacancies publicly. This obligation applies to it without qualification. Beaumont Ridge's job postings, whether on its own website or on third-party platforms, must state that AI is used in the screening, assessment or selection process where that is the case. The obligation extends to third-party recruiters screening on Beaumont Ridge's behalf. Non-compliant postings expose Beaumont Ridge to enforcement under the Employment Standards Act. No exemption applies here.

Court and tribunal duties apply where Beaumont Ridge is a party to proceedings. Canadian courts and several provincial tribunals have adopted or are developing requirements to disclose the use of AI in the preparation of filings, documents or evidence. Where Beaumont Ridge uses AI tools — including Microsoft 365 Copilot or ChatGPT — to draft litigation documents, regulatory submissions or tribunal filings, it must check the applicable court or tribunal rules in the relevant province for disclosure obligations. Failure to disclose AI use where required can result in adverse findings or sanctions. This is not a matter of best practice; it is a procedural obligation that must be verified at the outset of any proceeding.

The AI use register, PIA documentation, bias-testing records, consent records and publication of the privacy policy must all be retained for periods sufficient to respond to a regulatory investigation. PIPEDA does not prescribe a universal retention period, but the OPC expects records to be available for the duration of a complaint investigation. Quebec Law 25 requires that records supporting a PIA be retained for the life of the system to which they relate plus a reasonable period thereafter. Beaumont Ridge should adopt a minimum five-year retention standard for AI governance documentation.

6. What to Build Now Regardless of Parliamentary Timetable

Because AIDA will not be enacted and Bill C-36 has not yet received Royal Assent, there is no federal AI statute to anticipate in the short term. However, every building block Beaumont Ridge needs to be compliant with existing law also happens to be the foundation any future federal statute would require. Building now is not speculative preparation for hypothetical law; it is immediate compliance with PIPEDA, Quebec Law 25, Alberta PIPA, BC PIPA and O. Reg. 476/24, all of which are in force today.

The first priority is the AI use register and PIA programme. Beaumont Ridge should document every AI tool listed in the scope above, complete PIAs for the renewal-risk scoring model, the claims triage tool and the HR / recruitment systems within 60 days, and establish a schedule for reviewing those PIAs annually and on material change. The PIAs for tools that transfer data outside Canada — ChatGPT / OpenAI in particular — must assess adequacy of protection and document the contractual and technical safeguards in place.

The second priority is the section 12.1 notification and review mechanism for Quebec policyholders. Beaumont Ridge's renewal-risk scoring model, operating without broker involvement on straight cases, triggers this obligation now. Beaumont Ridge must implement a notification to Quebec policyholders at or before the point at which an automated renewal or pricing decision is communicated to them, explaining that the decision was made by automated processing, what personal information was used, and how they may request human review. A documented human review process with a named decision-maker must be operational before that notification goes live.

The third priority is job posting compliance under O. Reg. 476/24. All current public job postings must be reviewed and updated to include the required AI disclosure where AI is used. A standard disclosure clause must be inserted into the template for all future postings. This is a low-cost, immediate fix that eliminates a concrete enforcement risk.

The fourth priority is bias governance for the renewal-risk scoring model. Beaumont Ridge must commission a structured bias audit covering protected characteristics under the human rights codes of Ontario, Quebec, Alberta and British Columbia. The audit must use real or representative policyholder data, must be documented, and must produce a remediation plan where differential outcomes are identified. The audit should be repeated annually.

The fifth priority is vendor due diligence. Beaumont Ridge must obtain and retain data processing agreements with each AI vendor — Microsoft (Copilot), OpenAI (ChatGPT), the CRM vendor and the HR system vendor — that confirm the vendor processes personal information only on Beaumont Ridge's instructions, implements appropriate security measures, and does not use Beaumont Ridge's data to train its models without consent. Agreements that do not meet this standard must be renegotiated or the tool must not be used to process personal information.

The sixth priority is the French-language privacy policy. Beaumont Ridge must publish an updated privacy policy in both English and French that describes all AI uses, reflects the section 12.1 rights of Quebec residents, and provides contact details for the accountable individual. This is required under Quebec Law 25 and should be live without delay.

These six actions are not contingent on any future legislation. They address obligations that are in force today and that the OPC, the Commission d'accès à l'information du Québec and provincial privacy commissioners can enforce now. Completing them also positions Beaumont Ridge to demonstrate to insurers, commercial clients and regulators that its AI governance is structured, documented and maintained — a standard that will only become more important as OPC investigations and regulatory scrutiny of insurance-sector AI continue to intensify.

Federal and Provincial Compliance Matrix

One grid across both levels of government, so a control can be seen to satisfy several regimes at once.

Obligation	Strictest Regime (Default)	Federal — PIPEDA / Bill C-36	Quebec — Law 25	Ontario — PIPEDA + O. Reg. 476/24	Alberta — PIPA	British Columbia — PIPA	Effective Date / Sequencing
Lawful basis & consent for personal data processed by AI (health and financial data of policy holders)	Quebec Law 25 — explicit consent required for sensitive data; CAD $25m or 4% global turnover maximum fine	PIPEDA Sch. 1 Principles 4.3–4.4: meaningful consent; sensitive data requires explicit consent. Bill C-36 tabled 15 Jun 2026 — not yet in force.	Law 25 Arts. 12–14: explicit consent for sensitive data; s.12 requires purpose stated clearly. Highest penalty regime.	PIPEDA applies; no Ontario-specific private-sector privacy statute. Consent rules mirror federal baseline.	AB PIPA ss. 7–11: consent required; sensitive data needs explicit consent.	BC PIPA ss. 6–10: same consent framework; explicit consent for sensitive personal information.	Quebec Law 25 in force Sep 2023. PIPEDA ongoing. Bill C-36 — position to be confirmed on enactment.
Notification of automated decision-making (renewal-risk scoring model prices quotes with no broker — affects policy holders)	Quebec Law 25 Art. 12.1 — notification required where decision made exclusively by automated processing; strictest in Canada	No standalone federal automated-decision law. OPC guidance (2023) recommends transparency. Bill C-36 not yet in force.	Law 25 Art. 12.1: must inform individual before or at time of decision; must offer right to have decision reviewed by a person.	No Ontario-specific automated-decision obligation for private sector beyond PIPEDA/OPC guidance.	No AB PIPA automated-decision notification obligation beyond general accountability principle.	No BC PIPA automated-decision notification obligation beyond general accountability principle.	Quebec Law 25 Art. 12.1 in force Sep 2023. Beaumont Ridge must operationalise pre-decision notices for all Quebec policyholders immediately.
Right to human review of automated decisions (renewal pricing model — no broker involved on straight cases)	Quebec Law 25 Art. 12.1 — right to request human review of exclusively automated decisions; enforceable now	OPC guidance recommends meaningful human review for high-impact AI decisions; not yet a statutory right federally.	Law 25 Art. 12.1: individual may request that the decision be reviewed by a person. Beaumont Ridge must log and action such requests.	Not a statutory right in Ontario for private sector; best-practice standard applies via OPC guidance.	Not a statutory right under AB PIPA; accountability principle applies.	Not a statutory right under BC PIPA; accountability principle applies.	In force Sep 2023 (Quebec). Implement a documented review-request workflow for the renewal and quote model.
Obligation	Strictest Regime (Default)	Federal — PIPEDA / Bill C-36	Quebec — Law 25	Ontario — PIPEDA + O. Reg. 476/24	Alberta — PIPA	British Columbia — PIPA	Effective Date / Sequencing
AI bias and discrimination risk in automated underwriting and claims triage	Quebec Law 25 + Canadian Human Rights Act s.7 — combined	Canadian Human Rights Act s.7 prohibits discriminatory practices. PIPEDA	Law 25 Art. 12.1 review right creates indirect bias remedy. Commission des droits de	Ontario Human Rights Code s.9 prohibits discrimination in services and	AB Human Rights Act s.4 applies to services; automated outputs generating di	BC Human Rights Code s.8 applies. Claims triage routing that systematically disadvantage	All regimes in force. Annual bias audit of renewal-risk and claims-triage models

Obligation	Strictest Regime (Default)	Federal — PIPEDA / Bill C-36	Quebec — Law 25	Ontario — PIPEDA + O. Reg. 476/24	Alberta — PIPA	British Columbia — PIPA	Effective Date / Sequencing
(renewal-risk scoring; claims routing tool)	strictest exposure; dis criminatory algorithmic outcomes engage both regimes	accuracy principle (Sch. 1 s.4.6) requires data used in decisions to be accurate.	la personne oversight applies.	employment. Applies to automated underwriting output.	scriminatory outcomes engage this.	s protected groups is actionable.	required as of this compliance pack date.
Privacy impact assessment (PIA) for high-risk AI processing (health and financial data in renewal model, CRM, HR systems)	Quebec Law 25 Art. 3.3 — mandatory privacy impact assessment before any new technology project involving personal data; strictest obligation	PIPEDA Sch. 1 s.4.1.4: accountability includes PIA as best practice. OPC AI guidance (2023) treats PIA as expected for high-impact AI.	Law 25 Art. 3.3: PIA mandatory before imple mentation; must be com municated to Commission d'accès à l'information (CAI) on request.	No Ontario statutory PIA mandate for private sector; OPC guidance applies.	AB PIPA acc ountability principle supports PIA; not expressly mandatory by statute.	BC PIPA acc ountability principle supports PIA; not expressly mandatory by statute.	Quebec Law 25 in force Sep 2023. PIAs for all six in-scope AI tools must be completed and retained.
Breach notification to regulator and individuals where AI-processed personal data is compromised	Quebec Law 25 — 72-hour notification to CAI if risk of serious injury; strictest timeline in Canada	PIPEDA s.10.1: report to OPC and notify individuals of breaches posing real risk of significant harm. No fixed deadline but must be 'as soon as feasible'.	Law 25 Art. 3.5: notify CAI within 72 hours of becoming aware. Individual notification required where risk of serious injury. Record all incidents.	PIPEDA federal regime applies in Ontario. No Ontario-speci fic private-sector breach law adds to this.	AB PIPA s.34.1: notify OPC and individuals; no fixed hour deadline — 'without unre asonable delay'.	BC PIPA s.34.1: same as Alberta — notify without unreasonabl e delay.	Quebec 72-hour clock is operative now. Breach-response playbook must reflect the strictest deadline across all four provinces.
Accountabilit y and AI governance d ocumentation (policies, logs, model inventories for all six AI tools)	Quebec Law 25 Art. 3.3 + CAI enforcement practice — most prescriptive d ocumentation requirements	PIPEDA Sch. 1 s.4.1: designated Privacy Officer; maintain policies; OPC may audit. OPC AI guidance expects model inventories for high-impact tools.	Law 25: Privacy Officer mandatory (Art. 3.1); governance framework, register of processing activities, PIA register all required.	PIPEDA Privacy Officer requirement applies. Ontario has no additional private-sector governance mandate.	AB PIPA s.5: designated contact; policies required. Less prescriptive than Quebec.	BC PIPA s.5: same accoun tability structure as Alberta.	In force across all regimes. Beaumont Ridge must maintain a single consolidated AI and data governance register covering all provinces.

Obligation	Strictest Regime (Default)	Federal — PIPEDA / Bill C-36	Quebec — Law 25	Ontario — PIPEDA + O. Reg. 476/24	Alberta — PIPA	British Columbia — PIPA	Effective Date / Sequencing
Disclosure of AI use in job postings (Beaumont Ridge posts vacancies publicly; AI HR/recruitment system in use)	Ontario O. Reg. 476/24 — mandatory disclosure in every public job posting; in force 1 Jan 2026; applies to employers of 25+	No federal statute requires AI disclosure in job postings. OPC guidance encourages transparency in HR AI use.	No Quebec equivalent of O. Reg. 476/24 for private-sector job postings. Law 25 consent obligations apply to candidate data.	O. Reg. 476/24 under the Employment Standards Act: every publicly advertised posting must state whether AI is used to screen, assess or select. Applies to Beaumont Ridge (51–250 employees). Includes third-party screening.	No AB PIPA or employment statute equivalent to O. Reg. 476/24.	No BC PIPA or employment statute equivalent to O. Reg. 476/24.	Ontario O. Reg. 476/24 in force 1 Jan 2026. All Ontario job postings must already carry the disclosure. Review and update posting templates immediately.
Consent and transparency for AI use in performance management and HR decisions (AI HR/recruitment system)	Quebec Law 25 Art. 12.1 — automated decisions about employees or candidates require notification and review right; strictest	PIPEDA Sch. 1 s.4.3: consent for collection and use; employees have limited but real rights. OPC guidance applies to HR AI.	Law 25 Art. 12.1: if AI exclusively determines HR outcomes (shortlisting, performance scoring), notify employees and offer human review.	PIPEDA applies. Ontario ESA O. Reg. 476/24 covers candidate-facing disclosure; internal performance AI governed by PIPEDA and OPC guidance.	AB PIPA ss. 7–11 consent obligations apply to employee personal information. Workplace AI use must be disclosed in privacy notices.	BC PIPA ss. 6–10 apply to employee data. Privacy notice must describe AI-driven HR processing.	Quebec obligation in force Sep 2023. Ontario posting obligation in force 1 Jan 2026. Update employee and candidate privacy notices across all provinces.
Obligation	**Strictest Regime (Default)**	**Federal — PIPEDA / Bill C-36**	**Quebec — Law 25**	**Ontario — PIPEDA + O. Reg. 476/24**	**Alberta — PIPA**	**British Columbia — PIPA**	**Effective Date / Sequencing**
French-language obligations for AI-facing policyholder communications (Quebec clients; chatbot and customer service AI)	Quebec Charter of the French Language (Bill 96 amendments) — services to Quebec consumers must be available in French; strictest	No federal French-language obligation for private-sector AI customer interfaces beyond OLCA for federal bodies.	Charter of the French Language ss. 5, 52 (as amended): AI chatbot and automated correspondence for Quebec policyholders must be in French. Contracts must be in French	No Ontario equivalent mandatory French-language requirement for private-sector AI interfaces.	No AB equivalent.	No BC equivalent.	Charter of the French Language amendments (Bill 96) in force Jun 2022 and Jun 2023. Chatbot, policy documents, and renewal notices served to Quebec

Obligation	Strictest Regime (Default)	Federal — PIPEDA / Bill C-36	Quebec — Law 25	Ontario — PIPEDA + O. Reg. 476/24	Alberta — PIPA	British Columbia — PIPA	Effective Date / Sequencing
			unless client explicitly requests English.				clients must be French-first.
Client confidentiality controls when using generative AI tools (Microsoft 365 Copilot; ChatGPT/OpenAI processing client insurance data)	PIPEDA + Quebec Law 25 combined — transfer of sensitive client data to third-party AI provider requires contractual safeguards and may require consent	PIPEDA Sch. 1 s.4.1.3: accountable for data transferred to third parties; written contract required. OPC AI guidance: assess vendor data-retention and training practices.	Law 25 Art. 17: privacy impact assessment for cross-border or third-party data transfers; contractual protections mandatory. CAI may request PIA.	PIPEDA applies. No additional Ontario rule, but broker regulatory obligations under FSRA reinforce confidentiality duty.	AB PIPA s.13: disclosure to service providers permitted with contractual controls; no onward use for vendor's own purposes.	BC PIPA s.15.3: same service-provider transfer rule as Alberta.	All in force. Beaumont Ridge must review data-processing agreements with Microsoft and OpenAI; confirm no client data used to train external models; document findings in PIA register.
Declaration of AI use in regulatory or court filings (insurance regulator submissions; tribunal or court proceedings)	Court and tribunal rules — jurisdiction-specific; duty to declare AI-assisted drafting in filings before Ontario, Quebec, Alberta and BC courts and insurance tribunals	No federal statute mandates AI declaration in civil filings; Federal Court practice directions should be monitored.	Quebec Superior Court and administrative tribunals increasingly expect disclosure of AI-generated or AI-assisted content. Verify current practice directions.	Ontario courts — check current practice directions. Insurance appeals before Licence Appeal Tribunal (LAT): confirm LAT's AI declaration requirements.	Alberta courts and the Alberta Insurance Council: confirm current AI-disclosure practice directions for filings.	BC courts including BCSC: confirm current practice directions. BCFSA filings: check guidance.	Position should be confirmed per court or tribunal as rules evolve. Establish a pre-filing AI-use checklist for all provincial regulatory submissions.
OPC AI investigations and enforcement exposure (OPC actively investigating AI privacy matters as at Aug 2026)	OPC / Federal PIPEDA — OPC actively conducting AI investigations; findings create compliance benchmarks across all provinces	OPC has issued AI and privacy guidance; active investigations ongoing as at Aug 2026. Findings are public and set de facto standards. Cooperate fully with any OPC inquiry.	CAI is the Quebec enforcer; maximum fine CAD $25m or 4% of global turnover. CAI has audit powers under Law 25.	OPC is primary enforcer in Ontario for private sector. FSRA may refer data-protection concerns to OPC.	Office of the Information and Privacy Commissioner of Alberta enforces AB PIPA.	Office of the Information and Privacy Commissioner of BC enforces BC PIPA.	OPC investigations ongoing. Beaumont Ridge should designate an OPC liaison and maintain a regulatory-engagement log.

Obligation	Strictest Regime (Default)	Federal — PIPEDA / Bill C-36	Quebec — Law 25	Ontario — PIPEDA + O. Reg. 476/24	Alberta — PIPA	British Columbia — PIPA	Effective Date / Sequencing
Federal AI legislation — AIDA / Bill C-27 / Bill C-36 status (no assumptions; current position only)	Not yet a compliance obligation — AIDA definitively discontinued; Bill C-36 tabled 15 Jun 2026 but not in force	AIDA (part of Bill C-27) lapsed Jan 2025 and will not be reintroduced. Bill C-36 (priv ate-sector privacy reform) tabled 15 Jun 2026 — not yet enacted. No standalone federal AI statute is in force as at 11 Aug 2026.	Not applicable — federal bill does not affect Quebec Law 25 obligations, which are already in force.	Not applicable to Ontario compliance obligations currently in force.	Not applicable to AB PIPA obligations currently in force.	Not applicable to BC PIPA obligations currently in force.	Bill C-36 status to be confirmed on enactment. Monitor parliamentary progress. No action required under C-36 until Royal Assent.

RATING SCALE — cells use plain obligation statements. Where a regime does not apply, the cell states why. Where a regime applies but is less strict than the default, the cell records the applicable rule without the 'strictest' flag. The column 'Strictest Regime (Default)' identifies the single regime imposing the heaviest obligation or penalty for each row; controls designed to meet that regime will satisfy all lighter obligations in the same row unless a province imposes a structurally different requirement (e.g. Ontario O. Reg. 476/24 on job postings, which has no equivalent elsewhere). | MAXIMUM FINE NOTE — Quebec Law 25 maximum penalty is CAD $25,000,000 or 4% of worldwide turnover for the preceding fiscal year, whichever is greater. The turnover-based limb cannot be computed without Beaumont Ridge Insurance Brokers' revenue figures, which have not been provided. The absolute cap of CAD $25m is used throughout this matrix as the stated maximum. | CONTROL MAPPING — a single control satisfies multiple columns where marked. Example: a compliant consent notice drafted to Law 25 Art. 12.1 standard (automated-decision notification, human review right, bilingual French/English) simultaneously satisfies PIPEDA Sch. 1 ss. 4.3–4.4 (federal), AB PIPA ss. 7–11 and BC PIPA ss. 6–10. A single PIA template meeting Law 25 Art. 3.3 satisfies PIPEDA accountability guidance and AB/BC PIPA accountability principles. | REVIEW CADENCE — this matrix must be reviewed: (a) quarterly for the first year given active OPC investigations and Bill C-36 parliamentary progress; (b) immediately upon any Royal Assent or new OPC/CAI enforcement decision; (c) before each new AI tool deployment. Owner: Privacy Officer (to be designated under Law 25 Art. 3.1 and PIPEDA Sch. 1 s.4.1). | AI TOOLS IN SCOPE — all six tools identified in the compliance pack are covered: automated renewal-risk scoring and quote-pricing model; claims triage routing tool; AI-powered CRM/sales tools; Microsoft 365 Copilot; ChatGPT/OpenAI; AI in HR/recruitment systems. The renewal-risk and quote-pricing model is the highest-risk tool because it makes decisions affecting policyholders with no broker involvement on straight cases, directly engaging Quebec Law 25 Art. 12.1. | COURT AND TRIBUNAL FILINGS — this matrix records the obligation to declare AI use in filings but cannot specify practice directions for each court and tribunal as those are evolving. Beaumont Ridge must conduct a one-time review of current practice directions for: Quebec Superior Court, Court of Quebec, Ontario Superior Court, Licence Appeal Tribunal (Ontario), Alberta King's Bench, BC Supreme Court, and each relevant insurance regulatory body in all four licensed provinces. A pre-filing AI-use checklist must be implemented immediately. | BILL C-36 — tabled 15 June 2026; not in force as at 11 August 2026. Position should be confirmed on enactment. Do not treat C-36 obligations as live until Royal Assent. | NATIONAL AI STRATEGY — Canada's AI for All strategy launched 4 June 2026 is a policy framework, not directly enforceable legislation. Monitor for downstream regulatory instruments.

HR AI Procedures and Disclosure Templates

The full HR-side procedure set for AI in hiring, management and exit, with the templates to run it.

1. Scope and Governing Framework

These procedures govern all use of artificial intelligence tools by Beaumont Ridge Insurance Brokers in human resources processes, including hiring, performance management, discipline, and exit. They apply to all offices (Toronto and Montreal) and to all employees and candidates across Ontario, Quebec, Alberta, and British Columbia.

The legal framework governing these procedures includes: PIPEDA (federal), with Bill C-36 tabled 15 June 2026 but not yet in force and therefore not yet creating binding obligations; Quebec Law 25, enforced since September 2023, including its automated decision-making notification obligation; BC PIPA and Alberta PIPA; Ontario ESA and O. Reg. 476/24, in force 1 January 2026, requiring disclosure of AI use in publicly advertised job postings for employers of 25 or more; and OPC guidance on AI and privacy. AIDA (Bill C-27) lapsed January 2025 and will not be reintroduced; it creates no binding obligations but its risk-based principles inform Beaumont Ridge's internal standards.

AI tools in scope for HR purposes are: the AI in HR/recruitment systems (third-party platform), Microsoft 365 Copilot (used for drafting and knowledge management), and ChatGPT/OpenAI (used for internal drafting support). The automated renewal-risk scoring model and claims triage tool are out of scope for this document but are addressed separately in the compliance pack.

2. Sourcing Stage — Procedure and Permitted AI Use

At the sourcing stage, Beaumont Ridge may use AI tools to draft job descriptions, identify sourcing channels, and generate candidate outreach text. Microsoft 365 Copilot and ChatGPT may be used to draft vacancy copy. The HR/recruitment AI platform may be used to suggest sourcing channels based on historical hiring data.

AI may not autonomously publish a vacancy, set compensation bands without human approval, or make any decision to exclude a sourcing channel based on demographic inference. A designated HR manager must review and approve all vacancy text before publication.

All publicly advertised postings must include the Ontario ESA / O. Reg. 476/24 disclosure (see Template A below) because Beaumont Ridge employs more than 25 people and posts vacancies publicly. This obligation applies regardless of whether the posting is for an Ontario-based role, because the regulation applies to the employer, not the posting location, and Beaumont Ridge's head office subjects it to Ontario ESA coverage. Legal advice should be obtained if Beaumont Ridge wishes to argue a narrower scope, but the conservative and compliant position is to include the disclosure on all public postings.

Records retained at sourcing stage: copies of approved job descriptions, the version of the AI-generated draft before human editing, the name of the HR manager who approved the final text, and the date of publication. Retention period: three years from the date the vacancy closes, to align with limitation periods under provincial human rights legislation.

3. Screening Stage — Procedure and Permitted AI Use

At the screening stage, the HR/recruitment AI platform may rank or score applications against role criteria defined by HR. Microsoft 365 Copilot may be used to summarise CVs presented to a recruiter. ChatGPT may not be used to score or rank candidates because it is a general-purpose tool without the audit controls required for consequential HR decisions.

AI may assist by filtering out applications that do not meet objectively verifiable minimum criteria (e.g., possession of a required licence). AI may not make a final decision to reject a candidate. Every candidate rejected at screen must have their application reviewed by a human recruiter before rejection is confirmed. Where the AI platform flags a candidate as below threshold, the recruiter must view the flagged application and confirm or override the decision. Overrides must be logged with a reason.

Under Ontario ESA / O. Reg. 476/24, because Beaumont Ridge uses AI to screen applicants and employs 25 or more people, the job posting disclosure obligation is triggered. The disclosure must appear in the posting itself, not merely in a privacy notice. This requirement has been in force since 1 January 2026.

Under Quebec Law 25, where a decision is made exclusively by automated processing and has legal or significant effects on a person, Beaumont Ridge must notify the individual and, on request, allow them to have a human reconsider the decision. To avoid triggering this obligation inadvertently, no screen-stage rejection may be confirmed without human review, as stated above. If the process is ever changed so that the AI alone determines rejection, the Quebec notification obligation will apply immediately and the Accommodation and Alternative-Process route in Section 8 must be activated.

Records retained at screening stage: the AI platform's score or ranking output for each candidate, the recruiter's review notation, any override and its stated reason, and the final screen decision. Retention period: three years from vacancy close date.

4. Interview Stage — Procedure and Permitted AI Use

At the interview stage, Beaumont Ridge may use Microsoft 365 Copilot to generate structured interview question sets from a role profile approved by HR. The HR/recruitment AI platform may be used to schedule interviews and send calendar invitations. AI may not conduct, transcribe, or score interviews without explicit candidate consent, and even with consent, any AI-generated interview assessment must be reviewed by the hiring manager before it influences a decision.

AI-assisted video interview analysis tools are not currently listed in Beaumont Ridge's declared AI tools. If such a tool is introduced, a separate Privacy Impact Assessment (PIA) must be completed under PIPEDA and, for Quebec candidates, under Law 25 (which requires a PIA before implementing a technology that processes personal information), before deployment. This document does not authorise the use of any undeclared tool.

Interviewers must not use AI to generate a candidate assessment score based on facial expression, tone of voice, or other biometric inference. Such use would constitute processing of sensitive personal data without a valid basis under PIPEDA and would create material human rights exposure under provincial codes in Ontario, Quebec, Alberta, and BC.

Candidates who require accommodation during the interview process (e.g., alternative format, interpreter for French-language interviews in Quebec) must be offered an accommodation route as described in Section 8. Beaumont Ridge's obligation to serve clients and candidates in French under Quebec's Charter of the French Language applies to recruitment communications directed at Quebec candidates.

Records retained at interview stage: approved question set, names of interviewers, date and format of interview, any AI-generated scheduling or summary output, and the interviewer's independent written assessment. Retention period: three years from vacancy close date.

5. Selection Stage — Procedure and Permitted AI Use

At the selection stage, AI tools may provide comparative summaries of shortlisted candidates to assist the hiring manager. Microsoft 365 Copilot may compile a summary from notes and CVs. The HR/recruitment AI platform may present ranked scores.

AI may not make the selection decision. The hiring manager must make an independent, documented decision, stating the primary reasons for selection in writing. Where the AI platform's ranking and the hiring manager's decision differ, the hiring manager's decision governs and the discrepancy must be noted in the record.

Under Quebec Law 25, if selection were ever made exclusively by automated means, Beaumont Ridge would be required to notify the candidate of that fact, inform them of the personal information used, and provide them with the opportunity to submit observations to a human decision-maker who can reconsider the outcome. Beaumont Ridge's current process avoids this trigger by requiring a human selection decision in all cases. This procedure must not be changed without a legal review and update to these procedures.

Successful candidates receive an offer letter. Unsuccessful candidates at final shortlist stage receive a rejection communication. Neither communication may attribute the decision to the AI tool. The reason given to candidates should reflect the human decision-maker's stated reasoning.

Records retained at selection stage: AI comparative summary (if generated), hiring manager's written decision rationale, offer letter, rejection communications. Retention period: three years from vacancy close date, or longer if a human rights complaint is filed, in which case records must be preserved until the complaint is fully resolved.

6. Performance Management Stage — Procedure and Permitted AI Use

Microsoft 365 Copilot and the AI-powered CRM/sales tools may generate performance data summaries, call analytics, or activity metrics that feed into performance reviews for brokers and client-facing staff. These outputs are permitted inputs to a performance conversation but may not constitute the sole or primary basis for a performance rating, a performance improvement plan (PIP), or any formal performance outcome.

The line manager must conduct a substantive review of the AI-generated metrics, consider contextual factors (e.g., market conditions, client complexity, leave periods), and produce an independent written assessment. The AI output and the manager's assessment must both be retained.

Where AI tools are used to monitor employee productivity or communications (for example, Microsoft 365 Copilot's activity analytics), employees must be informed of this monitoring before it commences. Under PIPEDA, the collection of employee personal information must be for a purpose a reasonable person would consider appropriate, and employees must be notified of the purpose. Under Quebec Law 25, employees have rights to access personal information held about them and to have inaccurate information corrected. Under Alberta PIPA and BC PIPA, similar notification and access rights apply.

AI performance tools must not generate assessments based on protected grounds under applicable human rights legislation (including race, sex, disability, age, or family status). Any metric that could serve as a proxy for a protected ground (e.g., communication style scoring that disadvantages non-native English speakers) must be identified in the bias-testing process described in Section 9 and either adjusted or excluded.

Records retained at performance stage: AI-generated metric outputs, manager's written assessment, signed performance review form, any PIP documentation. Retention period: the duration of employment plus three years, consistent with Ontario ESA record-keeping requirements and equivalent provincial standards.

7. Discipline and Exit Stages — Procedure and Permitted AI Use

At the discipline stage, AI tools may be used to retrieve and organise factual records (e.g., call logs, CRM activity records, attendance data) relevant to a disciplinary matter. Microsoft 365 Copilot may assist HR in drafting disciplinary letters from a template approved by HR leadership. AI may not determine whether a disciplinary threshold has been met, recommend a disciplinary outcome, or generate a termination decision.

Every disciplinary decision must be made by a human manager with HR support. The reasons for the decision must be documented in writing and must not reference the AI tool as the basis for the outcome. Where AI-retrieved records are used as evidence, the employee must be provided with those records and given the opportunity to respond before a decision is made, consistent with procedural fairness requirements and applicable employment standards.

At the exit stage, AI tools may be used to generate off-boarding checklists and draft standard-form separation documentation. AI may not calculate termination entitlements under the Ontario ESA, Quebec Act Respecting Labour Standards, Alberta Employment Standards Code, or BC Employment Standards Act; those calculations must be completed by HR or legal counsel using verified statutory inputs, because errors create direct financial and legal liability.

Where an employee exits and was subject to AI-based performance monitoring, they retain the right under PIPEDA and applicable provincial privacy legislation to request access to personal information held about them, including AI-generated records, following their exit. Beaumont Ridge must be able to respond to such requests within the statutory timeframes (30 days under PIPEDA, with extension where permitted).

Records retained at discipline stage: all records used as evidence, meeting notes, decision letters, and any appeal outcome. Records retained at exit stage: separation agreement or termination letter, off-boarding checklist completion record, final pay calculation worksheet. Retention period for both: six years from date of exit, to cover the limitation period for wrongful dismissal claims in Ontario and equivalent periods in other provinces, and to meet Quebec's three-year limitation period with a conservative buffer.

8. Accommodation and Alternative-Process Route

Any candidate or employee who believes that an AI-assisted process has produced an unfair outcome, or who requires an alternative process because of a protected characteristic (including disability, language, religion, or family status), may invoke the accommodation route at any stage.

To invoke the route, the individual notifies HR in writing (email is sufficient) stating that they are requesting an accommodation or an alternative process, and the reason at a high level. Beaumont Ridge does not require a medical certificate or detailed personal disclosure at the point of initial request; further information may be sought only to the extent necessary to provide the accommodation.

Upon receipt of a request, HR must: (a) acknowledge receipt within two business days; (b) suspend any AI-generated output that has not yet resulted in a confirmed human decision; (c) assign a human reviewer who was not involved in the AI-generated output to assess the relevant stage afresh; and (d) communicate the outcome of the human review to the individual within five business days of acknowledgement, or notify them of an extended timeline if the matter requires more time.

For Quebec candidates and employees specifically: under Quebec Law 25, where a decision is made exclusively by automated means and has significant effects, the individual has a statutory right to be informed of that fact, to be told what personal information was used, and to submit observations to a human decision-maker who can reconsider the outcome. Beaumont Ridge's general procedure above satisfies this obligation provided it is followed in full. HR must not decline a reconsideration request from a Quebec individual on the grounds that the AI output was merely an input; the right attaches wherever the individual reasonably believes the AI output was determinative.

For candidates requiring French-language service in Quebec: all accommodation communications, notices, and decisions must be available in French. Template materials in this document must be translated and maintained in French by HR before use with Quebec candidates or employees.

Records retained: all accommodation requests, HR acknowledgements, human-review outputs, and outcome communications. Retention period: three years from the date of the request, or until any related complaint or litigation is resolved, whichever is later.

9. Bias Testing — Before and After Deployment

Before any AI tool is deployed or materially updated for use in HR processes at Beaumont Ridge, HR must complete a pre-deployment bias assessment. This assessment must: (a) identify all protected grounds under applicable human rights legislation that could be affected by the tool's outputs; (b) test the tool's outputs on a representative dataset for disparate impact across gender, age, ethnicity, and disability status at a minimum; (c) document the results and any remediation steps taken; and (d) obtain sign-off from the HR Director and a member of senior leadership before deployment proceeds.

For the HR/recruitment AI platform currently in use, a retrospective baseline bias assessment must be completed within 60 days of these procedures being adopted, if one has not already been conducted. The results must be retained.

Post-deployment, bias monitoring must be conducted at least annually and whenever a significant change occurs in the candidate or employee population, the role types for which AI is used, or the AI tool itself (including vendor-supplied model updates). Annual monitoring must review: selection rates at screen and selection stages disaggregated by gender, age group, and, where data is available and voluntarily provided, ethnicity; performance rating distributions disaggregated by the same categories; and any patterns in disciplinary outcomes linked to AI-flagged performance data.

Where the annual review identifies a statistically significant disparity that cannot be explained by a legitimate, job-related factor, Beaumont Ridge must: suspend the AI output for the affected stage pending investigation; conduct a root-cause review within 30 days; implement remediation (which may include retraining the tool, adjusting criteria, or discontinuing the AI input for that stage); and document the remediation and its outcome.

Bias testing reports must be retained for the life of the tool plus three years. They are subject to disclosure to the OPC, the Commission d'accès à l'information (Quebec), or any human rights tribunal upon request and may also be required in employment litigation. HR must treat these reports as confidential but must not suppress or alter an adverse finding.

10. Candidate Notice Template (Template A — Job Posting Disclosure, Ontario ESA / O. Reg. 476/24)

The following disclosure must appear in every publicly advertised job posting issued by Beaumont Ridge Insurance Brokers from 1 January 2026 onward. It must be legible and not buried in a hyperlink or footnote.

AI USE IN RECRUITMENT

Beaumont Ridge Insurance Brokers uses artificial intelligence tools to assist with the screening and assessment of applications for this role. AI may be used to rank or score applications against the requirements of this position. All AI-assisted outputs are reviewed by a human recruiter or hiring manager before any decision is made. If you have questions about this process or wish to request an alternative assessment method, please contact hr@beaumontridge.ca.

This disclosure satisfies O. Reg. 476/24 because it states that AI is used to screen and assess applicants. It is supplemented by the Privacy Notice in Template B. The French version of this disclosure (Template A-FR) must be used for postings directed at Quebec candidates or published on Quebec-facing platforms.

11. Candidate Privacy Notice Template (Template B — Collection of Personal Information)

The following notice must be provided to all candidates at the point of application, either embedded in the application form or provided as a clearly linked document.

PRIVACY NOTICE — RECRUITMENT

Beaumont Ridge Insurance Brokers collects personal information from you during the recruitment process for the purpose of assessing your suitability for the role you have applied for, communicating with you about your application, and, if you are successful, onboarding you as an employee.

Your personal information may be processed using AI-assisted tools, including an AI-powered recruitment platform and Microsoft 365 Copilot. These tools assist human reviewers; they do not make final decisions about your application without human involvement.

We may share your information with our HR software providers who act as service providers and are contractually required to protect your information. Your information will not be sold or used for any purpose unrelated to your application.

You have the right to access the personal information we hold about you and to request correction of inaccurate information. Quebec applicants have additional rights under Quebec Law 25, including the right to be informed of and to contest any decision made exclusively by automated means. To exercise your rights, contact privacy@beaumontridge.ca.

Application records are retained for three years following the close of the vacancy. If you are hired, your information becomes part of your employment record and is subject to our Employee Privacy Notice.

For questions about this notice, contact our Privacy Officer at privacy@beaumontridge.ca.

This notice must be translated into French and provided in French to all Quebec applicants (Template B-FR).

12. Employee AI Monitoring Notice Template (Template C — Existing Employees)

The following notice must be issued to all existing employees before any AI-based monitoring or performance analytics tool is activated or materially expanded.

NOTICE TO EMPLOYEES — USE OF AI IN PERFORMANCE MANAGEMENT

Beaumont Ridge Insurance Brokers uses AI-assisted tools as part of our HR and performance management processes. These tools include Microsoft 365 Copilot and an AI-powered CRM and sales analytics platform.

These tools may generate activity summaries, performance metrics, or data compilations that your manager may review as part of your performance discussions. They do not determine your performance rating, place you on a performance improvement plan, or result in any disciplinary outcome without a human manager reviewing the information and making an independent decision.

The personal information collected through these tools is used solely for lawful employment purposes and is protected in accordance with applicable privacy legislation, including PIPEDA and provincial privacy laws. You have the right to access personal information held about you and to request correction of inaccurate records.

If you believe an AI-generated output has produced an unfair result, or if you require an alternative process for any reason, you may request a human review by contacting HR at hr@beaumontridge.ca. Your request will be acknowledged within two business days.

For questions, contact our Privacy Officer at privacy@beaumontridge.ca.

This notice must be provided in French to employees based in Quebec or whose primary working language is French (Template C-FR). Evidence of delivery (e.g., email read receipt or signed acknowledgement) must be retained for the duration of employment plus three years.

13. Records Retention Summary

The following retention periods consolidate the requirements set out at each stage above and must be implemented in Beaumont Ridge's document management system.

Sourcing stage records (approved job descriptions, AI drafts, approval records): three years from vacancy close date.

Screening stage records (AI scores/rankings, recruiter review notations, overrides, screen decisions): three years from vacancy close date.

Interview stage records (question sets, interviewer assessments, scheduling records): three years from vacancy close date.

Selection stage records (AI comparative summaries, hiring manager decision rationale, offer and rejection letters): three years from vacancy close date, extended until resolution of any human rights complaint.

Performance management records (AI metric outputs, manager assessments, signed review forms, PIPs): duration of employment plus three years.

Discipline records (evidence records, meeting notes, decision letters, appeal outcomes): six years from exit date.

Exit records (separation documentation, off-boarding checklists, pay calculation worksheets): six years from exit date.

Accommodation records (requests, acknowledgements, human-review outputs, outcome communications): three years from the date of the request or resolution of related proceedings, whichever is later.

Bias testing reports: life of the tool plus three years.

Employee AI monitoring notices (delivery evidence): duration of employment plus three years.

All records that are subject to an active complaint, regulatory investigation, or litigation hold must be preserved beyond these standard periods until the matter is finally resolved. The Privacy Officer is responsible for issuing and lifting litigation holds. Quebec records involving personal information are also subject to the right of access and correction under Law 25; retention schedules must not result in destruction of records while an access request is pending.

Data Subject Rights Procedures for AI-Processed Data

How you answer an access, correction or deletion request when the data has passed through an AI system.

1. Purpose and Scope

This procedure governs how Beaumont Ridge Insurance Brokers ('Beaumont Ridge') handles requests from individuals to access, correct, or delete personal data that has passed through, or been processed by, any AI system operated by or on behalf of the firm. It applies to all six AI tools in scope: the automated renewal-risk scoring and quote-pricing model, the claims triage routing tool, the AI-powered CRM and sales tools, Microsoft 365 Copilot, ChatGPT/OpenAI integrations, and AI in HR and recruitment systems.

Beaumont Ridge holds health and financial information on policyholders, processes personal data in connection with commercial and personal lines insurance, and operates from offices in Toronto and Montreal with licences across Ontario, Quebec, Alberta, and British Columbia. The applicable legal framework includes PIPEDA (federal), Quebec Law 25 (Act respecting the protection of personal information in the private sector, as amended), Alberta PIPA, BC PIPA, and, for employment-related AI disclosures, Ontario ESA O. Reg. 476/24. All staff who receive or handle data subject requests must follow this procedure without deviation.

2. Rights, Statutory Deadlines, and Internal Owners

Right of Access. Under PIPEDA Principle 9, individuals may request access to their personal information and be told how it is or has been used. Beaumont Ridge must respond within 30 days of receiving a written request, extendable by a further 30 days where the volume or complexity of the request reasonably requires it, provided the individual is notified of the extension before the initial deadline expires. Under Quebec Law 25 (section 37 of the Act), the same 30-day rule applies. Alberta PIPA section 24 and BC PIPA section 23 also impose a 30-business-day response period. The internal owner for all access requests is the Chief Privacy Officer ('CPO'), supported by the Data & Systems Manager for locating AI-held data.

Right of Correction. PIPEDA Principle 9 and Quebec Law 25 (section 40) require Beaumont Ridge to correct inaccurate personal information and, where correction is declined, to annotate the record with the individual's statement of disagreement. Alberta PIPA section 26 and BC PIPA section 24 impose equivalent obligations. The 30-day response deadline applies in all jurisdictions. The internal owner is the CPO; the relevant business-unit manager (underwriting, claims, HR, or CRM) is responsible for implementing the correction in operational systems.

Right to Deletion. PIPEDA does not grant an absolute right of erasure, but requires Beaumont Ridge to destroy or anonymise personal information that is no longer necessary for the identified purposes (Principle 5). Quebec Law 25 (section 28) extends this to a positive obligation to destroy information once its purpose has been fulfilled, and the individual may request that destruction be accelerated where the firm has no lawful basis to retain it. Alberta PIPA section 35 and BC PIPA section 34 impose analogous retention-limitation duties. The CPO owns deletion requests; the Data & Systems Manager coordinates purge instructions to vendors.

Right to Explanation of an Automated Decision. Quebec Law 25 (section 12.1) requires that, where a decision is made exclusively by automated processing and produces legal or significant effects, the individual must be informed and may request that the decision be reviewed by a human and that reasons be given. This right is directly engaged by Beaumont Ridge's renewal-risk scoring and quote-pricing model (which handles straight cases without broker involvement) and the claims triage routing tool. The right must be communicated at the time of the decision. The CPO and the relevant Head of Underwriting or Head of Claims jointly own explanation requests.

Right to Human Review. Under Quebec Law 25 (section 12.1), any individual subject to an exclusively automated decision may request human review. Beaumont Ridge must make that review available and must confirm in writing that a qualified broker or claims handler has reviewed the case. No statutory deadline for the review itself is specified in the Act, but Beaumont Ridge adopts a 30-day internal target. PIPEDA's fairness principle and OPC guidance on AI require that individuals not be subjected to decisions based solely on automated processing without an opportunity to challenge them. The Head of Underwriting owns human-review requests for pricing/renewal; the Head of Claims owns them for triage decisions.

Right to Withdraw Consent / Object. Under PIPEDA Principle 3 and Quebec Law 25, individuals may withdraw consent to uses of their personal information, subject to legal or contractual restrictions. Where an individual objects to being subject to AI-driven profiling or scoring, Beaumont Ridge must assess whether a lawful basis other than consent justifies continued processing, and must respond in writing within 30 days. The CPO owns objection requests.

3. Locating Personal Data Held Inside or by an AI Vendor

When a data subject request is received, the Data & Systems Manager must conduct a structured data-mapping exercise across all six AI systems before any response is issued. This exercise uses Beaumont Ridge's AI Data Map (maintained as a living register by the CPO) and must be completed within ten business days of request receipt to leave adequate time for drafting the response.

For the renewal-risk scoring and quote-pricing model, the Data & Systems Manager must query the model's input database and output logs to identify all records linked to the individual, including raw policyholder attributes submitted for scoring and the resulting risk score and quoted premium. For the claims triage routing tool, the relevant claims-management system logs must be interrogated to locate all routing decisions and the data fields that drove them.

For Microsoft 365 Copilot, the Data & Systems Manager must use the Microsoft Purview compliance portal to run a content search across Exchange, Teams, SharePoint, and OneDrive for data associated with the individual's identifiers. Microsoft's data processing agreement confirms that Copilot prompt and response data is not used to train foundation models and is retained only as configured by the tenant administrator; Beaumont Ridge's retention policy must be verified against any active request.

For ChatGPT/OpenAI integrations, the Data & Systems Manager must review what personal data was submitted in prompts, using any API logging enabled by Beaumont Ridge. OpenAI's enterprise data processing terms confirm that API inputs are not used for model training by default; however, the CPO must verify that the current contract reflects this and that no opt-in to training has been activated. If personal data was submitted without API logging, this gap must be noted in the request log and disclosed to the individual.

For the AI-powered CRM and sales tools, the vendor's admin console must be queried for all contact, interaction, and scoring records associated with the individual. For AI in HR and recruitment systems, the HR Manager must extract all application data, AI-generated scores or assessments, and any structured or unstructured notes linked to the individual.

Where a vendor cannot confirm whether personal data has been incorporated into a model's weights through training, the CPO must note this uncertainty in the response, advise the individual of the limitation, and take the steps described in section 6 below.

4. Explaining an Automated Decision in Understandable Terms

Beaumont Ridge's renewal-risk scoring and quote-pricing model and claims triage routing tool both produce decisions or materially shape decisions affecting policyholders and claimants. Quebec Law 25 (section 12.1) requires that, upon request, the individual receive a plain-language explanation of the decision.

When an individual requests an explanation, the CPO must obtain from the Head of Underwriting or Head of Claims a written account that covers: (a) the categories of personal data that were inputs to the model (for example, claims history, property characteristics, credit-related attributes, or health information where applicable); (b) the principal factors the model identified as influencing the outcome in the individual's case; (c) the nature of the output (a risk score bracket, a price, or a triage category); and (d) the consequence of that output for the individual (for example, non-renewal, a higher premium, or routing to a specific claims pathway).

The explanation must avoid technical jargon and must be provided in the language in which the individual communicates with Beaumont Ridge. Where a policyholder is served in French, as is common in the Montreal office and across Quebec, the explanation must be provided in French in full compliance with the Charter of the French Language. The explanation must not merely reproduce a model's feature-importance output; it must be translated into a narrative that a non-specialist can act upon.

For HR and recruitment AI, where a candidate has been screened, assessed, or not selected in whole or in part through an AI tool, the explanation must describe the criteria the AI applied and the weight given to them. Ontario ESA O. Reg. 476/24 (in force 1 January 2026) requires public job postings to state whether AI is used to screen, assess, or select applicants; Beaumont Ridge employs more than 25 persons and posts vacancies publicly, so this disclosure obligation is already active. An explanation request from a candidate is therefore an extension of that existing transparency duty.

5. Right to Object and to Obtain Human Review of an Automated Decision

Any individual who has been subject to a decision made exclusively by automated processing — meaning the renewal-risk model operating on straight cases without broker involvement, or the claims triage tool routing a claim without human assessment — may request human review. This right is grounded in Quebec Law 25 (section 12.1) and is consistent with OPC guidance on fairness in automated decision-making.

On receipt of such a request, the CPO must immediately suspend any adverse action flowing from the automated decision (for example, a non-renewal notice or a claim denial based solely on triage output) pending completion of the review. The CPO must assign the review to a qualified, senior broker for underwriting matters or a senior claims adjuster for claims matters. That individual must assess the case independently, without being anchored to the model's output, using all available information including any additional representations the policyholder or claimant wishes to submit.

The reviewer must produce a written record of their assessment, noting where they agree with or depart from the model's output and the reasons. The outcome of the human review must be communicated to the individual in writing within 30 days of the request. Where the human review produces a different outcome from the model, the operational system must be updated and the original automated decision must be annotated accordingly.

For HR decisions, where a candidate requests human review of an AI-generated screening or assessment outcome, the HR Manager and the relevant hiring manager must jointly conduct a manual review of the application. The candidate must be informed of the outcome and given a brief explanation. Beaumont Ridge does not rely exclusively on AI to make final hiring decisions; however, where AI materially contributed to a rejection at screening stage, the human review must be substantive rather than a formality.

An individual's objection to being subjected to AI-driven profiling for renewal risk or claims routing must be assessed by the CPO against Beaumont Ridge's lawful basis for processing. Where consent was the basis, withdrawal of consent is effective; where the processing is necessary for the performance of the insurance contract, the CPO must advise the individual in writing that the processing cannot be avoided without affecting the availability of cover, and must offer a manually underwritten alternative where operationally feasible.

6. Deletion Where a Vendor May Have Retained or Trained on the Input

A deletion request requires Beaumont Ridge to act on its own systems and to instruct its vendors, but the ability to achieve complete erasure depends on how each vendor processes data. The CPO must pursue deletion through the following steps for each AI system.

For the renewal-risk scoring and quote-pricing model and the claims triage tool, where these are proprietary Beaumont Ridge systems or hosted models operating solely on firm data, the Data & Systems Manager must delete or anonymise all input records and output logs associated with the individual, subject to any lawful retention basis (see section 7). Where a retention basis exists (for example, a regulatory requirement to retain insurance records), the data must be flagged as subject to a deletion hold and must be purged as soon as that basis expires.

For Microsoft 365 Copilot, the CPO must issue a deletion instruction through the Microsoft Purview portal, covering all identified content containing the individual's personal data. Microsoft's contractual commitments confirm that Copilot does not train on tenant data; accordingly, model-weight contamination is not a risk with this tool, and deletion of the content in the tenant environment is sufficient.

For ChatGPT/OpenAI integrations, the CPO must submit a deletion request to OpenAI through the applicable data subject rights mechanism under Beaumont Ridge's enterprise agreement. Where the enterprise API contract confirms no training on inputs, deletion of Beaumont Ridge's own prompt logs satisfies the obligation. Where any uncertainty exists about training use, the CPO must record this in the request log, disclose it to the individual, and escalate to OpenAI's data protection team in writing, retaining the response.

For the AI-powered CRM and the HR/recruitment AI, the CPO must issue written deletion instructions to the respective vendors under the data processing agreements in place, specifying the individual's identifiers and requiring written confirmation of deletion within 15 business days. Where a vendor confirms that data has been incorporated into model training and cannot be isolated or removed from model weights, the CPO must: (a) document this limitation; (b) inform the individual plainly that deletion from model weights is technically not feasible; (c) ensure all other copies are deleted; and (d) assess whether the residual risk requires additional safeguards or notification to the relevant privacy authority. This situation must be escalated to the CPO immediately and not resolved at a lower level.

7. Refusal Grounds and Recording a Refusal

Beaumont Ridge may decline to fulfil a request, in whole or in part, only on the following grounds, each of which must be applied narrowly and documented in writing.

Under PIPEDA Schedule 1, Principle 9, access may be refused where: the information is protected by solicitor-client privilege; disclosure would reveal personal information about a third party that cannot be severed; the information was collected in relation to an investigation of a breach of an agreement or contravention of a law; or disclosure would reasonably be expected to threaten the life or security of another person. Alberta PIPA section 24(2) and BC PIPA section 23(3) contain equivalent grounds. Quebec Law 25 does not provide materially broader grounds for refusal.

Deletion may be refused where Beaumont Ridge is required by law to retain the information — for example, under insurance regulatory requirements in Ontario, Quebec, Alberta, or British Columbia, which mandate retention of policy and claims records for defined periods. In such cases, the CPO must identify the specific statutory or regulatory obligation in the refusal notice.

Where a refusal is issued, the CPO must: (a) inform the individual in writing of the refusal within the applicable statutory deadline; (b) state the specific ground relied upon and the legal provision; (c) advise the individual of their right to complain to the OPC (for PIPEDA matters), the Commission d'accès à l'information ('CAI') in Quebec, the Office of the Information and Privacy Commissioner of Alberta, or the Office of the Information and Privacy Commissioner for British Columbia, as applicable; and (d) enter the refusal in the Request Log with all required fields completed.

Refusals must never be based on administrative inconvenience, the cost of retrieval, or the fact that the data is held by a vendor. If a vendor declines to provide data or confirm deletion, that is an operational failure to be escalated to the CPO and addressed through the vendor contract, not a ground for refusing the individual.

8. Request Log Template

Beaumont Ridge must maintain a centralised, auditable Request Log for all data subject rights requests involving AI-processed data. The Log is owned by the CPO and must be reviewed quarterly. Each entry must capture the following fields in full.

Request reference number (sequential, unique). Date request received. Channel through which request was received (email, post, in person, web form). Full name of requestor. Relationship to Beaumont Ridge (policyholder, claimant, insurance applicant, job applicant, employee, other). Jurisdiction of requestor (Ontario, Quebec, Alberta, BC, other). Language of request (English or French). Type of right(s) invoked (access; correction; deletion; explanation of automated decision; human review of automated decision; objection to automated processing; withdrawal of consent; combination). AI systems implicated (list each of the six in-scope tools where relevant). Date acknowledgement sent to requestor. Date data-mapping exercise completed by Data & Systems Manager. Summary of data located and systems queried. Decision (fulfilled in full; fulfilled in part; refused). If refused in part or in full: ground relied upon, legal provision cited, and name of CPO approver. If human review invoked: name and role of reviewer, date review completed, outcome. Date final response sent to requestor. Statutory deadline applicable and whether met (yes / no — if no, reason). Vendor instructions issued (yes / no / not applicable) — if yes, list vendors and dates of instruction. Vendor deletion confirmation received (yes / no / outstanding / not feasible — if not feasible, detail recorded). Any regulatory notification made (yes / no — if yes, authority, date, reference). Date record closed. Any follow-up actions outstanding.

The Log must be stored in a manner that is itself compliant with PIPEDA and applicable provincial privacy laws, with access restricted to the CPO, the Data & Systems Manager, and the firm's legal advisers. It must be made available to any privacy authority on request. Entries may not be deleted from the Log even where the underlying personal data has been purged.

AI Incident Response Procedure

What happens in the first hour after an AI-related incident, and who does it. The regulator clock starts before anyone has decided who is in charge.

1. Purpose and Scope

This AI Incident Response Procedure governs how Beaumont Ridge Insurance Brokers ('Beaumont Ridge') detects, contains, escalates and reports incidents arising from the use of artificial intelligence systems across its Toronto and Montreal offices. It applies to all staff, managers, contractors and third-party vendors who operate or support AI tools on behalf of Beaumont Ridge.

The AI systems in scope are: the automated renewal-risk scoring and quote-pricing model; the claims triage routing tool; AI-powered CRM and sales tools; Microsoft 365 Copilot; ChatGPT / OpenAI interfaces; and AI used in HR and recruitment processes. These tools process personal data — including health and financial information on policyholders — and in some cases make or materially influence decisions about individuals. That combination attracts obligations under PIPEDA, Quebec Law 25, BC PIPA, Alberta PIPA, and Ontario ESA / O. Reg. 476/24, as well as OPC guidance on AI and privacy. The regulator's clock starts at the moment an incident is discovered, not when it is formally declared. Speed of action in the first hour is therefore a legal obligation, not merely good practice.

2. What Counts as an AI Incident

An AI incident is any event, near-miss or suspected failure involving an in-scope AI tool that has caused, or could cause, harm to individuals, regulatory breach, operational disruption, or reputational damage. The category is intentionally broad at the detection stage; severity is assessed separately.

Examples specific to Beaumont Ridge include: (a) the renewal-risk scoring or quote-pricing model producing outputs that systematically disadvantage policyholders on the basis of a protected characteristic such as ethnicity, age, disability or language group — including French-language clients served under Quebec Law 25; (b) the claims triage tool misrouting a claim in a way that delays or denies coverage to which a policyholder is entitled; (c) personal data — including health or financial records — being exposed to an unintended recipient through Microsoft 365 Copilot or a ChatGPT prompt, whether by output, log retention or training-data ingestion; (d) an HR recruitment AI screening out applicants without the required disclosure under Ontario O. Reg. 476/24, or applying criteria that constitute adverse impact discrimination; (e) the CRM tool generating inaccurate client-facing communications or fabricated policy information; (f) a vendor breach affecting any third-party AI system that holds or processes Beaumont Ridge client data; (g) the automated pricing model operating without a required broker review on a case that falls outside its defined straight-case parameters; (h) a court or tribunal filing containing AI-generated content that has not been reviewed and declared as required. Any of these events, whether confirmed or only suspected, triggers this procedure immediately.

3. Severity Levels and Triggered Response

Beaumont Ridge classifies AI incidents on three severity levels, assessed within the first fifteen minutes of discovery.

Level 1 — Critical: The incident involves actual or highly probable unauthorised disclosure of personal data (particularly health or financial data), a discriminatory automated decision affecting one or more individuals, a pricing or claims decision made without required human oversight, a vendor breach with confirmed data exposure, or any event that triggers a statutory notification obligation. Level 1 requires immediate escalation to the AI Incident Commander and Chief Privacy Officer (CPO), suspension of the affected tool, preservation of all logs and outputs, and commencement of the regulator notification assessment within the first hour. The Incident Commander must be reachable within ten minutes; if not reachable, the deputy assumes command without delay.

Level 2 — Significant: The incident involves a suspected (but unconfirmed) data exposure, a model output that is materially inaccurate and has been sent to clients or used in a decision, an HR AI disclosure failure that has not yet caused identifiable harm, or a CRM tool generating content that creates regulatory or reputational risk. Level 2 requires escalation to the AI Incident Commander within thirty minutes, isolation or supervised suspension of the affected tool, evidence preservation, and a formal decision within four hours on whether the incident should be reclassified to Level 1.

Level 3 — Low: The incident is a near-miss, an internal anomaly caught before any external effect, or a process deviation with no identified individual impact. Level 3 is logged in the AI Incident Register, reviewed by the AI Operations Lead within 24 hours, and assessed at the next scheduled governance meeting. No immediate tool suspension is required, but the affected workflow is flagged for review.

4. First-Hour Actions: Contain, Preserve, Stop

The first sixty minutes after discovery are governed by three mandatory parallel actions: contain the incident, preserve evidence, and stop the tool where required. These actions are taken concurrently, not sequentially.

Contain: The staff member who discovers the incident immediately notifies their line manager and the AI Incident Commander (or deputy) using the designated incident hotline. They do not attempt to fix, delete or correct the AI output. If the incident involves client-facing content — for example, an erroneous quote, a misdirected claims communication, or a chatbot response containing personal data — they do not send further communications using that tool until the Commander authorises resumption. If the incident occurred through a vendor system (for example, the OpenAI interface or the CRM provider), the vendor escalation contact is notified in parallel with internal escalation.

Preserve evidence: The discovering staff member, or the first technical responder if different, takes screenshots or exports of the AI output, the prompt or input that generated it, the timestamp, the user account involved, and any data fields accessed. System logs are placed on legal hold immediately. No logs, outputs or configurations are altered, overwritten or deleted. Where Microsoft 365 Copilot or ChatGPT is involved, any available interaction history is exported to a secure, access-restricted location. This evidence may be required for regulatory submissions, litigation, or insurance purposes.

Stop the tool: For Level 1 incidents, the affected AI system is suspended from live use immediately upon confirmation of the incident. For the renewal-risk or quote-pricing model, this means reverting affected cases to manual broker assessment. For the claims triage tool, all routed claims in the affected batch are flagged for manual review. For HR AI tools, no further automated screening outputs are acted upon pending review. The decision to suspend is taken by the AI Incident Commander; if unavailable within ten minutes, the deputy takes this decision without waiting for further authorisation. Suspension is logged with timestamp and authorising name.

5. Named Roles and Deputies

Clear personal accountability is essential because regulatory clocks run from the moment of discovery, not from the moment a role is filled. Every named role has a designated deputy who holds the same authority in the absence of the primary.

AI Incident Commander: The Chief Operating Officer (COO) or, in their absence, the Head of Operations. Responsible for overall incident management, tool suspension decisions, internal communication, and authorising external notifications. Reachable via the incident hotline at all times during business hours; an out-of-hours contact number is maintained in the Incident Register.

Chief Privacy Officer (CPO): Responsible for the regulator notification assessment, liaison with the OPC and provincial privacy commissioners, and oversight of any individual notification process. Deputy: the Legal and

Compliance Manager.

AI Operations Lead: Responsible for technical evidence preservation, vendor escalation, log management and tool configuration review. Deputy: the senior IT administrator.

HR Compliance Lead: Activated for any incident involving the HR or recruitment AI system, including Ontario O. Reg. 476/24 disclosure failures. Deputy: the HR Manager.

Legal and Compliance Manager: Advises on statutory deadlines, drafts regulatory notifications, and manages any court or tribunal declaration obligations. Deputy: external counsel on retainer (contact details in the Incident Register).

All deputies hold written delegated authority and are briefed on this procedure at least annually. Role assignments and deputy details are reviewed quarterly and updated in the AI Incident Register.

6. Regulator Notification: Decision and Statutory Deadline

The CPO commences the notification assessment as soon as a Level 1 incident is declared, and no later than the end of the first hour. The assessment is documented in writing, even if the conclusion is that notification is not required.

Under PIPEDA, Beaumont Ridge must notify the Office of the Privacy Commissioner of Canada as soon as feasible after determining that a breach of security safeguards involving personal information creates a real risk of significant harm to an individual. Significant harm includes humiliation, damage to reputation or relationships, financial loss, identity theft, and negative effects on a credit record. Health and financial data held on policyholders are inherently high-sensitivity, and any confirmed exposure must be assessed against this threshold promptly. There is no fixed number of days in the PIPEDA breach notification provision, but 'as soon as feasible' has been interpreted by the OPC to mean without unreasonable delay; a working assumption of 72 hours from confirmation of a reportable breach is operationally appropriate, and the CPO must be in a position to file or explain any longer interval.

Under Quebec Law 25 (An Act to modernise legislative provisions as regards the protection of personal information), Beaumont Ridge must notify the Commission d'accès à l'information (CAI) of any confidentiality incident involving personal information that presents a risk of serious injury to the person concerned, as soon as it becomes aware. Quebec Law 25 also requires notification where a decision is made exclusively by automated processing — the renewal-risk scoring model falls squarely within this provision, and any failure or discriminatory output from that model must be assessed under this head as well as the breach notification head.

Under Alberta PIPA and BC PIPA, notification obligations apply to privacy breaches causing a real risk of significant harm, with notification to the relevant commissioner required as soon as reasonably possible.

Bill C-36 was tabled on 15 June 2026 and, as at 11 August 2026, has not received Royal Assent. AIDA was definitively discontinued when Bill C-27 lapsed in January 2025 and will not be reintroduced. Neither creates current obligations. The CPO must not defer action on existing obligations pending C-36.

The maximum penalty under Quebec Law 25 is CAD $25 million or 4% of global turnover, whichever is greater. The turnover-based limb cannot be computed without Beaumont Ridge's financial figures, but the absolute cap is CAD $25 million. This exposure underscores that deferring the notification assessment is itself a risk.

7. Notifying Affected Individuals

The obligation to notify affected individuals is separate from, and may run in parallel with, the obligation to notify regulators. The CPO determines whether individual notification is required as part of the same first-hour assessment.

Under PIPEDA, individuals must be notified as soon as feasible if a breach of security safeguards creates a real risk of significant harm to them. The notification must contain sufficient information to allow the individual to understand the significance of the breach and to take steps to reduce the risk of harm. For policyholders whose health or financial data has been exposed, the threshold is almost certain to be met, and notification must not be delayed for investigative convenience.

Under Quebec Law 25, affected individuals must be notified of a confidentiality incident as soon as possible when there is a risk of serious injury. Where the automated pricing or claims triage model has produced a decision that has already been communicated to a policyholder and that decision is found to be erroneous or discriminatory, the individual must also be informed and the decision reviewed. Quebec Law 25 specifically requires that where a decision is made exclusively by automated processing that significantly affects a person, that person has the right to be informed of the decision, to have it explained, and to have a human being review it. The CPO ensures these rights are actioned as part of the incident response.

For FR-language clients in Quebec, individual notifications must be provided in French as a matter of law (Charter of the French Language) and as a matter of Beaumont Ridge's operational obligation to its Montreal client base.

Alberta PIPA and BC PIPA similarly require individual notification where a breach creates a real risk of significant harm. In all cases, the notification records the date sent, the method, the content, and the name of the authorising officer. Copies are retained in the Incident Register. Where notification is assessed as not required, the documented reasoning is retained for a minimum of two years.

8. Vendor Escalation Route

Several of Beaumont Ridge's AI incidents will originate in, or be exacerbated by, third-party systems. The vendor escalation route is activated in parallel with internal response, not after it is complete.

For Microsoft 365 Copilot: escalation is through the Beaumont Ridge Microsoft account manager and, for confirmed data incidents, through the Microsoft Security Response Center (MSRC) portal. The AI Operations Lead holds the current service agreement and support contact details. Beaumont Ridge's data processing agreement with Microsoft must be reviewed immediately to confirm the contractual notification obligations Microsoft owes to Beaumont Ridge, and any Microsoft-side breach notification received is treated as a Level 1 trigger.

For ChatGPT / OpenAI: the AI Operations Lead contacts OpenAI through the enterprise support channel. Staff are reminded that inputting unredacted client personal data into ChatGPT outside a formally contracted and data-processed environment may itself constitute a breach; the incident assessment must determine whether the root cause includes unauthorised data submission.

For the AI-powered CRM and the HR recruitment system: the AI Operations Lead contacts the named vendor support contact documented in the Incident Register and requests a written statement of the scope of any data involved, within four hours of escalation. The vendor's own breach notification obligations do not discharge Beaumont Ridge's independent obligations to regulators and individuals.

For the renewal-risk scoring and claims triage tools, if these are operated by or with a third-party model provider, the same escalation protocol applies. Where a vendor fails to respond within the contractual support window, the Legal and Compliance Manager is notified immediately, as delayed vendor response does not pause the regulatory clock. All vendor communications in connection with an incident are preserved as evidence.

9. Ontario AI Disclosure — HR Recruitment Incidents

Ontario O. Reg. 476/24 under the Employment Standards Act came into force on 1 January 2026. It requires that any publicly advertised job posting made by an employer of 25 or more employees — including where a third party screens on their behalf — must state whether artificial intelligence is used to screen, assess or select applicants. Beaumont Ridge employs between 51 and 250 staff and posts vacancies publicly; this obligation applies in full.

An AI incident in the HR context includes: posting a vacancy without the required AI disclosure; using an HR AI tool to screen or reject applicants without that disclosure having been made; or the HR AI system producing outputs that reflect bias against a protected group under the Ontario Human Rights Code or applicable federal human rights legislation.

Where such an incident is identified, the HR Compliance Lead is immediately notified. Affected postings are corrected and reissued with the required disclosure. Where applicants may have been screened or rejected through an undisclosed AI process, the HR Compliance Lead determines, in consultation with the Legal and Compliance Manager, whether those applicants must be notified and whether their applications must be reconsidered. This determination is documented and recorded in the Incident Register. Employment Standards Officers may request evidence of compliance; the preserved posting records and screening logs are retained for a minimum of three years.

10. Court and Tribunal Obligations

Beaumont Ridge staff and external counsel acting on its behalf must declare the use of AI in preparing court or tribunal filings where the relevant court or tribunal rules so require. As at 11 August 2026, several Canadian courts and administrative tribunals have issued practice directions or standing orders requiring disclosure of AI-generated content in submissions.

An incident under this head arises where AI-generated content has been included in a filing without the required declaration, or where AI output has introduced factual inaccuracies into a submission. On discovery, the Legal and Compliance Manager is notified immediately. Corrective steps — which may include filing an amended document with a declaration and a correction notice — are taken without delay. The Incident Register records the filing reference, the AI tool used, the content involved, the corrective action taken, and the date. This category of incident does not typically engage privacy regulator notification, but it does engage professional regulatory obligations and potential adverse costs consequences; the Legal and Compliance Manager assesses each case on its facts.

11. Post-Incident Review and the AI Incident Register

Every incident, regardless of severity level, concludes with a post-incident review and a record in the AI Incident Register. The review is not optional and is not deferred pending resolution of any regulatory process.

For Level 1 incidents, the post-incident review is conducted within five business days of containment, led by the CPO and AI Incident Commander. It examines the root cause, the adequacy of the initial response, whether the affected tool's controls or parameters require amendment, whether vendor contractual terms require renegotiation, and whether this procedure itself requires updating. The review produces a written report with dated action items and named owners.

For Level 2 incidents, the review is conducted within ten business days, led by the AI Operations Lead with CPO input.

For Level 3 incidents, the review is conducted at the next scheduled AI governance meeting, which occurs at least quarterly.

The AI Incident Register is a permanent, access-controlled record maintained by the Legal and Compliance Manager. Each entry includes: date and time of discovery; discoverer's name and role; description of the incident; AI tool(s) involved; severity classification and rationale; first-hour actions taken with timestamps; regulator notification assessment outcome and, if notification was made, the date and reference; individual notification outcome; vendor escalation record; post-incident review date and report reference; and any follow-up action status. The Register is reviewed in full at each quarterly AI governance meeting and is available for production to the OPC, CAI, or any other competent regulator on request. Entries are retained for a minimum of five years. The Register also serves as the evidence base for Beaumont Ridge's annual AI transparency and accountability reporting under OPC guidance.