

Specimen — illustrative example

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

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

AI Risk Management Standard

Beacon Talent Partners

Document type	standard	Version	1.0 Draft
Owner	Head of Delivery	Customer decision role	Managing Director
Effective date	Upon customer release	Review cycle	Annual
Framework basis	RiskWoven AI practical best practices		
Jurisdiction	Ontario, Canada		
Prepared for	Beacon Talent Partners	Industry	Recruitment and staffing
Organization size	26-100	Locations	Not specified

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

1. Purpose and Scope

This standard sets the mandatory requirements Beacon Talent Partners must meet when assessing, testing, reviewing, and monitoring AI systems. It implements the commitments

made in the AI Governance and Accountability Policy and the AI Acceptable Use Policy.

This standard applies to every AI system that Beacon Talent Partners operates, procures, or uses to support recruitment and staffing activity, including the candidate ranking model and any AI-assisted candidate outreach drafting. It applies to all staff who configure, operate, or act on the output of those systems, and to any vendor whose AI system processes Beacon Talent Partners data.

This standard does not duplicate the step-by-step intake and review steps set out in the Third-Party AI Review Procedure. Where a system is vendor-supplied, that procedure runs alongside this standard.

2. Definitions

AI risk. The possibility that an AI system produces an outcome that causes harm to a candidate, client, or Beacon Talent Partners. Harm includes unfair treatment, privacy breach, reputational damage, regulatory exposure, and operational failure.

Impact. The severity of harm if a risk materialises. For this standard, impact is assessed across five categories: candidate harm (including unfair treatment or exclusion), privacy and data harm, operational harm, reputational harm, and legal or regulatory harm. Fairness is treated as a standalone impact category: any outcome that could produce a disparate result across a protected characteristic carries at least a High impact rating regardless of the other categories.

Likelihood. The probability that a risk materialises given the controls currently in place. Likelihood is rated Low, Medium, or High based on the frequency of the activity, the maturity of the control, and any evidence of prior failure.

Evaluation. A structured assessment of an AI system's performance, fairness, accuracy, and control effectiveness, conducted before the system is used and repeated on the schedule set in Section 8.

Drift. A change in an AI system's behaviour or output distribution over time, relative to its behaviour at the point of evaluation, that has not been caused by a deliberate and reviewed update. Drift may affect accuracy, fairness, or both.

Residual risk. The risk that remains after all controls have been applied. Residual risk is the basis for the review decision recorded in the AI Risk Register.

3. Document Ownership and Accountability

The Head of Delivery owns this standard. The Head of Delivery is responsible for maintaining it, ensuring evaluations are completed on schedule, and keeping the AI Risk Register current.

The Managing Director holds review authority for this standard and for any AI system assessed as High residual risk. The Managing Director's review record must be captured in the AI Risk Register before a High-risk system is used.

All staff who operate or act on AI system outputs are responsible for complying with the requirements in this standard and for reporting anomalies to the Head of Delivery without delay.

4. AI System Documentation

Before any AI system is evaluated under Section 5, the Head of Delivery must ensure a system record exists that covers all of the following.

4.1 Intended purpose. A plain-English statement of what the system is used for. For the candidate ranking model, this is: ordering applicants against a role specification to produce a shortlist for recruiter review, and drafting candidate outreach messages.

4.2 Affected people. The individuals whose data the system processes or whose outcomes the system influences. For the ranking model, affected people are candidates whose CVs are submitted against a role.

4.3 Benefits. The operational or quality improvements the system is expected to deliver.

4.4 Limitations. The conditions under which the system's output is less reliable, the data it cannot interpret correctly, and any known gaps in its training or design. The system record must note that candidate CVs unavoidably contain information from which protected characteristics can be inferred, and that the system's ability to exclude such signals is a known limitation requiring active monitoring.

4.5 Foreseeable misuse. The ways the system could be used outside its intended purpose, including use to infer age, health, family status, or any other protected characteristic. The system record must state that such use is prohibited under the AI Acceptable Use Policy.

4.6 Data inputs. The categories of data the system receives. For the ranking model: candidate CVs, employment history, education records, and role specifications. The system record must

note that candidate data may not be used to train the vendor's model and that all processing must remain within Canada, consistent with the Privacy and Personal Information Policy.

4.7 Decision output. What the system produces and how that output is used. The system record must state that the ranking model produces a ranked shortlist; it does not reject candidates automatically. Low-ranked candidates are less likely to be read by a recruiter, which means the ranking carries material consequence even without automatic rejection.

The system record is stored in the AI Risk Register. The Head of Delivery reviews it for accuracy whenever the system is re-evaluated.

5. Risk Assessment

Every AI system in scope must be assessed before use and reassessed on the schedule in Section 8. The Head of Delivery conducts the assessment. The assessment is recorded in the AI Risk Register.

5.1 Risk domains. The assessment must address each of the following domains.

Data risk. Whether the data inputs are accurate, complete, and appropriate for the system's purpose. Whether data handling complies with the Data Classification and Handling Standard and the Privacy and Personal Information Policy.

Privacy risk. Whether the system processes personal information in a way that could breach applicable Ontario privacy law. Whether candidates are informed that automated ranking is used and that they may request human review.

Security risk. Whether the system and its data connections meet the requirements of the Information Security Policy. Whether access to the ranking model's request logs and the applicant tracking system is restricted to authorised staff.

Fairness risk. Whether the system's output could produce a disparate outcome across a protected characteristic. This domain is assessed for every system in scope. Any feature that proxies for a protected characteristic must be identified and removed on discovery. Any risk in this domain is rated High impact regardless of likelihood.

Accuracy risk. Whether the system's output is sufficiently reliable for its intended use. Whether the system's performance has been validated against Beacon Talent Partners' own candidate population.

Transparency risk. Whether the basis for the system's output can be explained to a candidate who requests it. Whether the features used in each ranking decision are logged in a form that supports that explanation.

Human-oversight risk. Whether the human review controls are sufficient given the system's influence on candidate outcomes. For the ranking model, this includes whether recruiters read every shortlist, whether a sample of candidates below the shortlist cut is reviewed on each role, and whether override decisions are recorded.

5.2 Rating scale. Each identified risk is rated on two dimensions.

Impact: Low, Medium, or High, assessed across the five impact categories in Section 2. The highest category rating applies. Any fairness risk is rated High impact.

Likelihood: Low, Medium, or High.

Residual risk is the combination of impact and likelihood after controls are applied. The Head of Delivery records the residual risk rating in the AI Risk Register alongside the controls relied upon.

5.3 High-risk threshold. A system is classified as High residual risk if any single assessed risk carries a High impact rating after controls, or if the combination of Medium-rated risks indicates that a material adverse outcome is plausible. The candidate ranking model is treated as High residual risk by default because of its fairness risk profile and its influence on candidate outcomes.

5.4 Fairness as a mandatory category. The fairness domain is not optional and cannot be rated as not applicable. If the assessment cannot produce a fairness finding because the vendor has not supplied sufficient information about the model's features or training data, that gap is itself recorded as a High-risk finding.

6. Acceptance Criteria and Pre-Use Testing

A system may not be used until it has passed the acceptance criteria set out in this section and the review record required by Section 7 has been completed.

6.1 Acceptance criteria. The Head of Delivery sets acceptance criteria for each system before testing begins. Criteria must be specific enough that a reviewer can determine whether each one has been met. At a minimum, criteria must address:

- **Accuracy:** the system's output must meet a minimum performance threshold set by the Head of Delivery and recorded in the AI Risk Register before testing begins. The information supplied does not supply this threshold; the Head of Delivery must define it as an organisational parameter before the first evaluation is conducted.
- **Fairness:** the system must not produce a statistically significant difference in shortlist inclusion rates across candidate groups defined by protected characteristics, at a significance threshold set by the Head of Delivery and recorded in the AI Risk Register. The information supplied does not supply this threshold; the Head of Delivery must define it before the first evaluation.
- **Feature compliance:** no feature that proxies for a protected characteristic may be active in the model at the point of use.
- **Transparency:** the features used in each ranking decision must be logged in the ranking model request log in a form that supports explanation to a candidate.
- **Human oversight:** the recruiter override mechanism must be functional and override records must be written to the override log.
- **Data residency:** all processing must occur within Canada.
- **Vendor data use:** the vendor must have confirmed in writing that candidate data is not used to train the vendor's model.

6.2 Pre-use testing. Before a system is used for the first time, and before it is returned to use after a material change, the Head of Delivery must conduct or commission testing against each acceptance criterion. Testing must use a dataset that is representative of Beacon Talent Partners' candidate population. Test results are recorded in the AI Risk Register alongside the date of testing and the name of the person who conducted it.

6.3 Failed criteria. If any acceptance criterion is not met, the system must not be used until the criterion is met or an exception is recorded under Section 9. A failed criterion is not resolved by adjusting the criterion to match the result.

7. Review Record

7.1 High-risk systems. A system assessed as High residual risk may not be used until the Managing Director has reviewed the completed risk assessment, the test results, and the residual risk rating, and has recorded that review in the AI Risk Register. The review record must state the date of review, the reviewer's name, the residual risk rating reviewed, and any conditions attached to use.

7.2 Non-high-risk systems. A system assessed as Medium or Low residual risk may be reviewed by the Head of Delivery. The review record must state the same fields as 7.1.

7.3 Review conditions. Where a reviewer attaches conditions to use — such as a requirement to re-evaluate within a shorter period, or a restriction on the roles for which the system may be used — those conditions are recorded in the AI Risk Register and treated as mandatory controls. Breach of a review condition is treated as a control failure under Section 10.

8. Ongoing Monitoring

8.1 Monitoring scope. The Head of Delivery is responsible for ongoing monitoring of each AI system in use. Monitoring covers accuracy, fairness, feature compliance, human oversight, and data residency.

8.2 Log sources. Monitoring draws on the following sources:

- Applicant tracking system access logs
- Ranking model request logs, including the features used in each decision
- Recruiter override records
- Identity sign-in logs

Ranking and override records are retained for three years to support any challenge to a decision. Technical logs are retained for a period set by Beacon Talent Partners and recorded in the AI Risk Register; 90 days searchable and one year archived is offered as a starting point and is not derived from any cited control. All retention periods must be consistent with the Records Retention and Disposal Policy; Beacon Talent Partners must confirm that alignment before this standard takes effect.

8.3 Monitoring activities. The Head of Delivery must carry out the following monitoring activities on a cycle set by Beacon Talent Partners and recorded in the AI Risk Register.

Quarterly is offered as a starting point for each activity and is not derived from any cited control; Beacon Talent Partners must set the actual cycle before this standard takes effect.

- Review a sample of ranking decisions to check that shortlist composition is consistent with the fairness acceptance criteria.
- Review override records to confirm that recruiter overrides are being recorded and that the rate of override is within the range expected at the point of evaluation.
- Check that no new feature proxying for a protected characteristic has been introduced by a vendor update.
- Confirm that candidate data remains within Canada.
- Confirm that the vendor's written confirmation of no model training on candidate data remains current.

8.4 Drift detection. The Head of Delivery must compare current ranking output distributions against the baseline recorded at the most recent evaluation on a cycle set by Beacon Talent Partners; quarterly is offered as a starting point and is not derived from any cited control. If drift is detected — meaning the output distribution has shifted beyond a tolerance threshold set by the Head of Delivery and recorded in the AI Risk Register — the system must be re-evaluated under Section 5 before further use. The information supplied does not supply the drift tolerance threshold; the Head of Delivery must define it as an organisational parameter.

8.5 Candidate transparency. The Head of Delivery must confirm on each monitoring cycle that the candidate-facing disclosure stating that automated ranking is used, and the mechanism by which a candidate may request human review, remain accurate and accessible.

8.6 Monitoring records. Each monitoring activity is recorded in the AI Risk Register with the date, the activity performed, the finding, and any action taken. Overdue monitoring activities are reported to the Managing Director.

9. Exceptions

9.1 Requesting an exception. Where Beacon Talent Partners cannot meet a requirement of this standard, the Head of Delivery must record an exception in the AI Risk Register. The exception record must state: the requirement that cannot be met, the reason, the compensating control in place, the residual risk with the compensating control, the expiry date of the exception, and the name of the reviewer.

9.2 Review authority for exceptions. Any exception that relates to a High-risk system or a fairness control must be reviewed by the Managing Director. Other exceptions may be reviewed by the Head of Delivery.

9.3 Expiry. An exception expires on the date recorded in the exception record. An expired exception with no replacement is a control failure. The Head of Delivery must report the number of exceptions past their expiry date to the Managing Director on the monitoring cycle set under Section 8.3.

9.4 No standing exceptions. An exception may not be renewed indefinitely. If the underlying gap cannot be resolved within a period set by Beacon Talent Partners and recorded in the AI Risk Register, the Head of Delivery must escalate to the Managing Director for a decision on whether the system may continue in use.

10. Reassessment Triggers

The Head of Delivery must initiate a full reassessment under Section 5 when any of the following occurs.

- A material change to the system, including a vendor update that alters the model's features, training data, or output logic.
- An AI incident involving the system, as defined in the AI Incident Response Procedure.
- A privacy breach involving candidate data processed by the system, as defined in the Privacy Breach Response Procedure.
- Detection of drift beyond the tolerance threshold set under Section 8.4.
- A new legal or regulatory obligation that affects how the system may be used.
- A finding from monitoring that a fairness acceptance criterion is no longer being met.
- A complaint from a candidate about the system's output.

Reassessment follows the same process as the initial assessment. The system must not return to use after a material change or incident until the reassessment is complete and the review record under Section 7 has been updated.

11. Metrics and Reporting

The Head of Delivery reports the following metrics to the Managing Director on an annual basis, aligned with the annual review of this standard. Where a metric falls outside the

expected range, the Head of Delivery must report it to the Managing Director without waiting for the annual cycle.

- Number of High-risk systems reviewed by the Managing Director before use, as a proportion of all High-risk systems in use.
- Number of monitoring activities completed on schedule, as a proportion of all monitoring activities due in the period.
- Percentage of required reassessments completed within the period triggered by the events in Section 10.
- Number of exceptions past their expiry date at the point of reporting.
- Number of material control failures identified in the period and number of corrective actions overdue at the point of reporting.

These metrics are recorded in the AI Risk Register. The Head of Delivery is responsible for the accuracy of the data.

12. Relationship to Other Documents

This standard implements the AI Governance and Accountability Policy and the AI Acceptable Use Policy. It does not restate the obligations in those policies.

Data handling requirements are set in the Data Classification and Handling Standard. Privacy obligations are set in the Privacy and Personal Information Policy. Security requirements are set in the Information Security Policy. Third-party vendor requirements are set in the Third-Party Risk Management Policy. The step-by-step intake process for vendor AI systems is set in the Third-Party AI Review Procedure.

Incident response steps are set in the AI Incident Response Procedure, the Security Incident Response Procedure, and the Privacy Breach Response Procedure. This standard determines when those procedures are triggered; it does not duplicate their steps.

Log retention periods stated in Section 8.2 must be consistent with the Records Retention and Disposal Policy. Beacon Talent Partners must confirm that alignment before this standard takes effect.

13. Review and Maintenance

The Head of Delivery reviews this standard annually and submits any proposed changes to the Managing Director for review. The Managing Director's review record is captured in the document version history.

The Head of Delivery must also initiate an out-of-cycle review if a material change to Beacon Talent Partners' AI systems, a significant incident, or a new legal or regulatory obligation makes the current version of this standard inadequate.

The version history records the date of each review, the reviewer's name, the changes made, and the reason for any change.

Owner review items

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

Type	Detail	Where
gap	The information supplied does not specify the jurisdiction of Beacon Talent Partners' locations beyond Ontario, Canada. If staff or candidates are located in other provinces or countries, additional privacy and AI obligations may apply. Beacon Talent Partners' legal reviewers should confirm which instruments apply — candidates for review include PIPEDA, Ontario's privacy legislation, and any sector-specific regulation — but this standard names no statute because the supplied material does not confirm which instruments apply.	Section 5.1 — Privacy risk; Section 12
gap	The accuracy performance threshold and the fairness significance threshold required by Section 6.1 are not supplied. Both are mandatory organisational parameters that must be defined before the first evaluation is conducted. Without them, the acceptance criteria cannot be tested.	Section 6.1
gap	The monitoring cycle for Section 8.3 and 8.4 activities has not been set. Quarterly is offered as a starting point. Beacon Talent Partners must set the actual cycle before this standard takes effect.	Section 8.3; Section 8.4
gap	The drift tolerance threshold required by Section 8.4 is not supplied. This is a mandatory organisational parameter. Without it, drift detection cannot produce a testable finding.	Section 8.4
gap	The log retention periods in Section 8.2 are stated by Beacon Talent Partners in their answers. The technical log retention period is offered as a starting point and must be set by Beacon Talent Partners as an organisational parameter. All retention periods must be confirmed as consistent with the Records Retention and Disposal Policy before this standard	Section 8.2; Section 12

	takes effect. This standard cannot verify that alignment because the text of that policy was not supplied.	
gap	The maximum duration for an exception before mandatory escalation (Section 9.4) has not been set. Beacon Talent Partners must set this as an organisational parameter.	Section 9.4
vague	The candidate outreach drafting use case is noted in the system record requirement but is not separately risk-assessed in the supplied answers. If the drafting tool uses candidate data or produces outputs that could affect candidate experience, it should be assessed as a distinct system or as a distinct use case of the ranking model. The Head of Delivery should confirm whether a separate system record and assessment are required.	Section 4.1; Section 5
gap	The information supplied does not confirm whether the AWS environment used for the ranking service is subject to any contractual or regulatory data residency requirement beyond the organisation's own rule that data stays in Canada. Beacon Talent Partners' legal and procurement reviewers should confirm this, particularly in relation to applicable Ontario privacy law and the vendor agreement.	Section 5.1 — Data risk; Section 6.1 — Data residency
gap	The sibling document AI Governance and Accountability Policy is referenced as the parent policy for this standard. This standard cannot assert what that policy says because its text was not supplied. Beacon Talent Partners must confirm that the accountability structure, risk appetite, and definitions in this standard are consistent with that policy.	Section 1; Section 3

Information the organization must supply

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

- Accuracy performance threshold for the candidate ranking model (Section 6.1): the minimum acceptable performance level, expressed as a measurable metric, that the system must meet before use. This is an organisational parameter; the Head of Delivery must define the metric and the threshold value before the first evaluation is conducted.
- Fairness significance threshold (Section 6.1): the maximum acceptable difference in shortlist inclusion rates across candidate groups defined by protected characteristics, expressed as a statistical threshold. The Head of Delivery must define this before the first evaluation.
- Monitoring cycle for Section 8.3 and 8.4 activities: the actual frequency at which each monitoring activity must be completed. Quarterly is offered as a starting point and is not a requirement of any cited control.
- Drift tolerance threshold (Section 8.4): the maximum permitted shift in output distribution before a reassessment is triggered. The Head of Delivery must define this as a measurable value.

- Maximum exception duration before mandatory escalation (Section 9.4): the period after which an unresolved exception must be escalated to the Managing Director for a decision on continued use.
- Technical log retention period (Section 8.2): the actual searchable and archive retention durations to be set by Beacon Talent Partners as organisational parameters, and confirmation that all retention periods are consistent with the Records Retention and Disposal Policy.
- Confirmation of whether the candidate outreach drafting tool requires a separate system record and risk assessment, or whether it is covered as a use case within the ranking model's existing record.
- Confirmation from legal and procurement reviewers of which privacy and AI instruments apply to Beacon Talent Partners in Ontario, and whether any additional obligations arise from the AWS hosting arrangement or from candidates or clients located outside Ontario.
- Confirmation that the accountability structure, risk appetite, and definitions in this standard are consistent with the AI Governance and Accountability Policy.

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