

Specimen — illustrative example

Northvale Health Analytics 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

Northvale Health Analytics

Document type	standard	Version	1.0 Draft
Owner	Chief Technology Officer	Customer decision role	Chief Executive Officer
Effective date	Upon customer release	Review cycle	Annual
Framework basis	RiskWoven AI practical best practices		
Jurisdiction	Ontario, Canada		
Prepared for	Northvale Health Analytics	Industry	Health technology
Organization size	26-100	Locations	Not specified

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

1. Purpose

This standard sets the mandatory, testable requirements that Northvale Health Analytics applies when assessing, testing, reviewing, and monitoring AI systems. It implements the

mandates in the AI Governance and Accountability Policy and the AI Acceptable Use Policy. It applies alongside the Privacy and Personal Information Policy, the Information Security Policy, and the Data Classification and Handling Standard, which govern the underlying data and security controls. Where this standard and a sibling document address the same topic, the more specific requirement governs.

2. Scope

This standard applies to every AI system that Northvale Health Analytics operates, procures, or materially configures, including systems operated by third parties on Northvale's behalf. It applies to all staff, contractors, and vendors who build, deploy, configure, or use those systems. It covers the full lifecycle of each system: from initial evaluation through deployment, ongoing monitoring, and decommissioning.

Two AI systems are currently in scope:

- (a) Triage Assistant — reads inbound referral letters from family physicians and proposes a priority band and a suggested specialty routing. A registered clinician confirms or overrides every suggestion before it takes effect.
- (b) Referral History Summariser — produces a summary of long referral histories for the reviewing clinician. No clinical decision is made on the basis of the summary alone.

Any new AI system, or any material change to an existing one, falls within this standard from the point at which evaluation begins.

3. Definitions

AI risk — the possibility that an AI system produces an outcome that causes harm to patients, clinicians, the organisation, or the public. Harm includes clinical harm, privacy breach, discriminatory output, reputational damage, and regulatory non-compliance.

Impact — the severity of harm if a risk materialises. Northvale rates impact on a four-point scale: 1 (negligible), 2 (minor), 3 (significant), 4 (critical). Any outcome that touches patient safety is treated as critical regardless of the score a reviewer would otherwise assign.

Likelihood — the probability that a risk materialises given the controls in place. Northvale rates likelihood on a four-point scale: 1 (rare), 2 (unlikely), 3 (possible), 4 (likely).

Evaluation — the structured process of assessing an AI system's intended purpose, risks, controls, and fitness for use before it is deployed or materially changed.

Drift — a measurable change in an AI system's behaviour or output distribution over time, relative to its behaviour at the point of evaluation, that was not caused by a deliberate update.

Residual risk — the risk that remains after all controls have been applied. Residual risk is the basis on which a review decision is made and monitoring requirements are set.

4. Roles and Accountabilities

Chief Technology Officer (CTO) — owns this standard. The CTO ensures every AI system in scope has a completed risk record, that monitoring is performed on schedule, and that this standard is reviewed annually. The CTO reviews application access logs, model request and override logs, Azure administrative activity, and identity sign-in logs.

Chief Executive Officer (CEO) — holds review authority for high residual risk. The CEO records the review decision, the conditions attached to it, and the date by which the risk must be reassessed.

System Owner — named for each AI system in the AI system register. The system owner is accountable for the accuracy of the system's risk record, for executing monitoring tests, and for raising a reassessment when a trigger event occurs. Where no system owner is named, the CTO holds that accountability.

Registered Clinician — confirms or overrides every priority band proposed by the Triage Assistant before it takes effect. The clinician's confirmation or override is recorded in the model request and override log.

All Staff — must report suspected AI failures, unexpected outputs, or misuse to the CTO without delay, following the AI Incident Response Procedure.

5. AI System Register

The CTO maintains an AI system register. Every AI system in scope must have an entry before it is used in any live or production context.

Each register entry must record:

- (a) System name and version.
- (b) Intended purpose, stated in plain language.
- (c) The population affected and the nature of the effect.
- (d) The data categories processed, cross-referenced to the Data Classification and Handling Standard.
- (e) The system owner.
- (f) The current risk rating (likelihood × impact) and residual risk rating.
- (g) The review record: who reviewed, the date, and any conditions attached.
- (h) The monitoring schedule and the date monitoring was last completed.
- (i) The date the next reassessment is due.

The register is reviewed by the CTO on a cycle set by the CTO and recorded before this standard takes effect. Quarterly review is one suggested starting point for the CTO's consideration; it is not a default requirement and is not derived from any cited control. The register is the authoritative source for the status of every AI system.

6. Risk Assessment Requirements

6.1 When a risk assessment is required

A risk assessment must be completed before any AI system is deployed to a live environment. A new assessment must be completed before any material change takes effect. A reassessment must be completed when a trigger event occurs (see Section 10).

6.2 What the assessment must document

Every risk assessment must document all of the following:

- (a) Intended purpose — what the system is designed to do, the decisions or outputs it produces, and the clinical or operational context in which it operates.
- (b) Affected people — the patients, clinicians, or other individuals whose data is processed or whose outcomes may be influenced by the system's outputs.
- (c) Benefits — the specific, measurable improvements the system is expected to deliver.
- (d) Limitations — the conditions under which the system is known or expected to perform poorly, including data types it was not designed to handle and populations underrepresented in its training.

- (e) Foreseeable misuse — the ways in which the system could be used outside its intended purpose, including uses prohibited by the AI Acceptable Use Policy.
- (f) Data risk — whether the system processes personal health information, whether processing is confined to Azure Canada Central, whether identifiers beyond the referral itself are excluded, and whether patient data is used for training or fine-tuning. No patient data may be used to train or fine-tune any model. No data may leave Canada.
- (g) Privacy risk — assessed in alignment with the Privacy and Personal Information Policy. The assessment must identify the legal basis for processing and any obligations triggered under applicable Ontario privacy law.
- (h) Security risk — assessed in alignment with the Information Security Policy. The assessment must identify the threat vectors relevant to the system and the controls that address them.
- (i) Fairness risk — whether the system's outputs may disadvantage identifiable groups of patients or clinicians, and what evidence exists that the system performs consistently across those groups.
- (j) Accuracy risk — the error rate the system produces in conditions representative of live use, the consequences of each error type, and the threshold below which accuracy is unacceptable.
- (k) Transparency risk — whether the system's outputs are explainable to the clinician who acts on them, and whether the basis for a proposed priority band can be retrieved and recorded.
- (l) Human-oversight risk — whether the human review step is effective, whether a clinician can meaningfully evaluate the system's proposal, and whether override rates indicate that the system's proposals are being accepted without genuine review.

6.3 Risk rating

Each identified risk is rated by multiplying the likelihood score (1–4) by the impact score (1–4), producing a score of 1–16. Any risk that touches patient safety is escalated to the highest impact rating regardless of the score produced by the multiplication.

Risk bands:

- Low: 1–4
- Medium: 5–9
- High: 10–16, or any risk touching patient safety

6.4 Residual risk

After controls are identified and applied, the assessor re-rates each risk to produce a residual risk score. The residual risk score, not the inherent score, determines the review authority and monitoring requirements.

SPECIMEN

7. Acceptance Criteria and Pre-Deployment Testing

7.1 Acceptance criteria

Before a system is submitted for review, the system owner must define acceptance criteria for each of the following dimensions. Criteria must be specific and measurable so that a reviewer can determine whether they have been met.

- (a) Accuracy — the minimum acceptable performance metric (for example, precision, recall, or F1 score) on a test dataset representative of live referral letters. The specific threshold is set by the system owner and recorded in the risk record. The system owner must state the basis for the threshold.
- (b) Fairness — the maximum acceptable performance gap between patient subgroups, defined by the system owner and recorded in the risk record.
- (c) Override rate — the range of clinician override rates that indicates the system is functioning as intended. An override rate outside that range is a trigger for investigation. The system owner sets the acceptable range and records it.
- (d) Data boundary compliance — 100% of processing must occur within Azure Canada Central. No patient data may be transmitted outside Canada. This criterion admits no exception.
- (e) Human confirmation — 100% of priority band proposals must be confirmed by a registered clinician before taking effect. This criterion admits no exception.

7.2 Pre-deployment testing

The system owner must conduct and document testing against each acceptance criterion before submitting the system for review. Testing must use data that is representative of live conditions. Where live patient data is used in testing, it must be handled in accordance with the Privacy and Personal Information Policy and the Data Classification and Handling Standard.

The test record must state:

- The dataset used and its characteristics.
- The test conducted for each criterion.
- The result.

- Whether the criterion was met.
- The name of the person who conducted the test and the date.

A system that fails any acceptance criterion may not proceed to review until the failure is remediated and the test is repeated.

7.3 Third-party AI systems

Where the AI system is operated by a third party, the system owner must follow the Third-Party AI Review Procedure. Evidence of third-party testing must be obtained and assessed against Northvale's acceptance criteria before the system is reviewed for deployment.

8. Review Authority and Review Conditions

8.1 Review authority by residual risk

Low residual risk — reviewed by the CTO. The review is recorded in the AI system register.

Medium residual risk — reviewed by the CTO. The review record must include the conditions under which the system may operate and the date by which the risk must be reassessed.

High residual risk — reviewed by the CEO. The CEO's review record must state:

- The residual risk accepted.
- The conditions under which the system may operate.
- The fixed period for which the review applies.
- The date by which the risk must be reassessed.
- Any corrective actions required and their due dates.

No high-residual-risk system may operate beyond the period stated in the CEO's review record without a new review.

8.2 Review conditions

The reviewer may attach conditions to a review decision. Conditions are mandatory controls that must be in place for the review to remain valid. The system owner is accountable for maintaining compliance with all conditions. A condition that is no longer met triggers an immediate reassessment.

8.3 Prohibited uses

The following uses are prohibited regardless of risk rating or review status, and no review decision may authorise them:

- (a) The Triage Assistant assigning a priority band without clinician confirmation.
- (b) Any AI system communicating directly with a patient.
- (c) Any AI system being used to decline or close a referral.
- (d) Processing patient data outside Azure Canada Central.
- (e) Using patient data to train or fine-tune any model.
- (f) Transmitting any data outside Canada.

9. Monitoring Requirements

9.1 Ongoing monitoring

The system owner must monitor each deployed AI system on a schedule set by the CTO and recorded in the AI system register before deployment. The CTO must set and record the monitoring frequency for each system; monthly monitoring is one suggested starting point for the CTO's consideration for high-residual-risk systems and is not a default requirement and is not derived from any cited control.

Monitoring must cover:

- (a) Accuracy — the system's performance against the acceptance criteria defined at deployment. A result below the accepted threshold is a trigger for reassessment.
- (b) Drift — comparison of current output distribution against the baseline recorded at deployment. The system owner must define the drift threshold that triggers a reassessment and record it in the risk record.
- (c) Override rate — the rate at which clinicians override the Triage Assistant's proposals, reviewed against the acceptable range defined at deployment. An override rate outside that range triggers investigation.
- (d) Data boundary compliance — confirmation that all processing remains within Azure Canada Central and that no data has left Canada.
- (e) Human confirmation compliance — confirmation that 100% of priority band proposals were confirmed by a registered clinician before taking effect.

(f) Access and override logs — the CTO reviews application access logs, model request and override logs, Azure administrative activity, and identity sign-in logs on a cycle set by the CTO and recorded before deployment. The review is recorded. Weekly review of override logs is one suggested starting point for the CTO's consideration for the Triage Assistant; it is not a default requirement and is not derived from any cited control.

9.2 Log retention

Override and model-decision log retention periods are an organisational parameter that the CTO must set and record before deployment. Northvale's existing practice is to retain override and model-decision logs for 7 years alongside the clinical record; 7 years is noted here as the organisation's current practice and is not a requirement imposed by this standard. Technical log retention is similarly an organisational parameter; Northvale's existing practice is to keep technical logs searchable for 90 days and archived for 1 year. The Records Retention and Disposal Policy governs disposal.

9.3 Monitoring records

Each monitoring test must produce a written record stating the date, the person who conducted the test, the metric tested, the result, and whether the result was within the accepted range. Records are held by the CTO and are available for internal review.

9.4 Metrics reported to the CTO

The system owner reports the following metrics to the CTO on the monitoring schedule:

- Percentage of high-residual-risk systems reviewed before use.
- Percentage of scheduled monitoring tests completed on time.
- Percentage of required reassessments completed on time.
- Number of review decisions past their stated expiry date.
- Number of material control failures and overdue corrective actions.

10. Reassessment Triggers

The system owner must initiate a reassessment without delay when any of the following occurs:

(a) A material change to the AI system, including a model update, a change to the data inputs, a change to the integration with clinical workflows, or a change to the human review step.

- (b) An AI incident is raised under the AI Incident Response Procedure that involves the system.
- (c) A performance metric falls below an accepted threshold or outside an accepted range.
- (d) Drift is detected at or above the threshold defined in the risk record.
- (e) A new legal or regulatory obligation applies to the system or the data it processes.
- (f) A condition attached to a review decision is no longer met.
- (g) The review period stated in a high-residual-risk review record expires.
- (h) A clinician or staff member raises a concern about the system's outputs that cannot be resolved without re-evaluation.

The system owner notifies the CTO when a reassessment trigger occurs. The CTO determines whether the system must be suspended pending reassessment. A system that poses an immediate patient safety risk must be suspended until the reassessment is complete and a new review decision is recorded.

11. Annual Review of This Standard

The CTO reviews this standard annually. The review must assess whether the requirements remain appropriate given changes to Northvale's AI systems, applicable law, and the organisation's risk appetite. The review record is held by the CTO. Any material change to this standard requires a new review record signed by the CTO and, where the change affects high-residual-risk controls, reviewed by the CEO.

12. Document Control

Owner: Chief Technology Officer

Review frequency: Annual

Next review due: Set by the CTO at the point of adoption

Version history: Maintained by the CTO

This standard is advisory output produced by RiskWoven AI. It does not constitute professional advice, certification, accreditation, or a compliance determination. Northvale Health Analytics should obtain independent professional review before relying on this standard to satisfy any regulatory obligation.

Owner review items

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

Type	Detail	Where
gap	The organisation's jurisdiction is Ontario, Canada, and the systems process personal health information. The applicable privacy legislation (which may include PHIPA and PIPEDA or its successor) is not named in the supplied material. The body of this standard refers only to 'applicable Ontario privacy law' to avoid asserting statutory content that was not supplied. Northvale's professional reviewers must confirm which instruments apply and whether any notification, retention, or breach-reporting obligations affect the requirements in this standard.	Section 6.2(g)
gap	The supplied material does not specify the monitoring frequency for any AI system. This standard requires the CTO to set and record a frequency before deployment. The CTO must ensure this is done for both the Triage Assistant and the Referral History Summariser before this standard takes effect.	Section 9.1
gap	The supplied material does not specify the access and override log review frequency. This standard requires the CTO to set and record a review cycle before deployment.	Section 9.1(f)
gap	The supplied material does not specify the accuracy threshold, fairness gap threshold, or acceptable override rate range for either AI system. These are acceptance criteria that the system owner must define and record in the risk record before pre-deployment testing begins.	Section 7.1
gap	The supplied material does not specify the drift threshold for either AI system. The system owner must define and record this threshold in the risk record before deployment.	Section 9.1(b)
gap	The AI system register does not yet exist as a named document in Northvale's document set. The CTO must create and maintain it. This standard depends on it as the authoritative record for all AI systems.	Section 5
gap	The supplied material does not specify the AI system register review frequency. This standard notes quarterly as one suggested starting point for the CTO's consideration. The CTO must set and record the actual frequency.	Section 5
gap	The supplied material does not name a system owner for either the Triage Assistant or the Referral History Summariser. Until system owners are named, the CTO holds that accountability by default under this standard. The CTO should assign named system owners as soon as practicable.	Section 4
vague	The supplied material states that high residual risk requires CEO acceptance 'for a fixed period, recorded with a review date' but does not specify what that period is. This standard requires the CEO to state the period in the review record. The CTO and CEO should agree a default maximum period before the first high-residual-risk review is conducted.	Section 8.1
gap	The Third-Party AI Review Procedure is referenced in Section 7.3. This standard does not duplicate its requirements. The CTO must confirm that the procedure covers the evidence-gathering steps needed to satisfy Northvale's acceptance criteria for third-party systems.	Section 7.3

gap	The AI Governance and Accountability Policy and the AI Acceptable Use Policy are referenced as the policy instruments this standard implements. This standard has not been given the text of those policies. The CTO must confirm that the requirements in this standard are consistent with the mandates in both policies, particularly the prohibited uses listed in Section 8.3.	Sections 1 and 8.3
gap	The supplied material does not specify physical locations for Northvale's offices or staff. If staff access systems from locations outside Ontario or Canada, the data boundary requirements in Sections 6.2(f) and 8.3 may need additional controls. The CTO should confirm the geographic scope of staff access.	Section 2
gap	Log retention periods are organisational parameters that the CTO must set and record before deployment. The periods noted in Section 9.2 reflect Northvale's existing practice as supplied and are not requirements imposed by this standard. The CTO should confirm these periods are consistent with applicable legal obligations.	Section 9.2

Information the organization must supply

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

- Confirmation of which Ontario and federal privacy instruments apply to Northvale's processing of personal health information, obtained from independent professional review.
- Named system owners for the Triage Assistant and the Referral History Summariser.
- The monitoring frequency for each AI system, set by the CTO and recorded in the AI system register before deployment.
- The access and override log review frequency, set by the CTO and recorded before deployment.
- The accuracy threshold, fairness gap threshold, and acceptable override rate range for each AI system, set by the system owner and recorded in the risk record before pre-deployment testing.
- The drift threshold for each AI system, set by the system owner and recorded in the risk record before deployment.
- The AI system register, created and maintained by the CTO as the authoritative record for all AI systems in scope.
- The AI system register review frequency, set by the CTO.
- A default maximum period for high-residual-risk review decisions, agreed between the CTO and CEO before the first such review is conducted.

- Confirmation that the Third-Party AI Review Procedure covers the evidence-gathering steps needed to satisfy Northvale's acceptance criteria for third-party AI systems.
- Confirmation that the requirements in this standard are consistent with the mandates in the AI Governance and Accountability Policy and the AI Acceptable Use Policy.
- Confirmation of the geographic scope of staff access to AI systems, to determine whether additional data boundary controls are required.
- Confirmation of the log retention periods to be applied, and whether those periods satisfy any applicable legal or regulatory obligations.

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.

SPECIMEN