

Specimen — illustrative example

Meridian Community Credit Union 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 Governance and Accountability Policy

Meridian Community Credit Union

Document type	policy	Version	1.0 Draft
Owner	Chief Risk Officer	Customer decision role	Board Risk Committee
Effective date	Upon customer release	Review cycle	Annual
Framework basis	RiskWoven AI practical best practices		
Jurisdiction	Ontario, Canada		
Prepared for	Meridian Community Credit Union	Industry	Financial services
Organization size	101-500	Locations	Not specified

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

1. Purpose

This policy establishes Meridian Community Credit Union's governance framework for artificial intelligence systems. It sets out who is accountable for AI decisions, how AI systems are

inventoried and classified, what must happen before an AI system is deployed or materially changed, how systems are monitored throughout their operational life, and how they are retired. The policy exists to ensure that AI use at Meridian is proportionate to risk, transparent to members, and subject to meaningful human oversight.

2. Scope

This policy applies to every AI system that Meridian operates, procures, or materially configures, regardless of whether it runs on Meridian's own infrastructure or is delivered by a third party. It applies to all staff, contractors, and suppliers who develop, deploy, operate, or advise on AI systems on Meridian's behalf.

This policy does not govern the use of AI tools by individual staff for personal productivity. That is addressed in the AI Acceptable Use Policy. Supplier engagement obligations are addressed in the Third-Party Risk Management Policy and the Third-Party AI Review Procedure. Incident response steps are addressed in the AI Incident Response Procedure.

3. Definitions

AI system means any automated system that uses a model, algorithm, or learned function to generate outputs — including scores, recommendations, classifications, or decisions — that influence a business process or a member outcome. This includes the pre-adjudication lending model and any successor or additional system brought into scope under this policy.

System owner means the named Meridian employee who holds accountability for a specific AI system throughout its lifecycle: from pre-deployment assessment through to retirement. The system owner is responsible for ensuring the system operates within its approved parameters and that all required reviews, records, and escalations occur.

Risk tier means the classification assigned to an AI system based on the potential severity of harm to members, the credit union, or third parties if the system produces an incorrect, biased, or otherwise harmful output. Risk tiers are defined in the AI Risk Management Standard.

Impact assessment means the structured evaluation completed before an AI system is deployed or materially changed. It examines the system's purpose, the data it uses, the population it affects, the potential for harm including disparate impact on member groups, and the controls that will govern it. The assessment produces a risk tier assignment and identifies the conditions under which the system may operate.

Material change means any modification to an AI system that could alter its outputs, its risk profile, or the population it affects. This includes changes to the underlying model or its training data, changes to the input data sources, changes to the decision thresholds or output categories, changes to the human oversight arrangements, and changes to the supplier or hosting environment. The system owner determines whether a proposed change is material, and that determination is recorded.

4. Policy Statement

Meridian Community Credit Union uses AI systems only where they serve a clear member or business purpose, where the risks are understood and proportionate, and where appropriate human oversight is in place. No AI system may operate at Meridian without a current risk tier assignment, a named system owner, and a record of the review that permitted its deployment.

Meridian treats fairness and the potential for disparate impact on member groups as a distinct risk category, not as a subset of operational risk. This applies to every AI system that influences a member outcome.

Members have the right to request human review of any AI-influenced decision affecting them. That right is unconditional and is not subject to any threshold, fee, or prior condition.

5. Governance Structure and Accountability

Board Risk Committee holds ultimate oversight of AI governance at Meridian. It reviews and accepts high residual risk AI use cases, reviews the annual AI governance report, and sets the risk appetite within which AI systems operate. Critical supplier relationships — including the core banking provider, the scoring model vendor, and the credit bureau — require Board Risk Committee review before engagement or material change.

Chief Risk Officer (CRO) owns this policy and is accountable for the AI governance framework. The CRO ensures the AI system inventory is maintained, that impact assessments are completed and reviewed, that the risk tier methodology is applied consistently, and that escalations reach the Board Risk Committee when required. The CRO presents the annual AI governance report to the Board Risk Committee.

System owners are accountable for the day-to-day governance of their assigned AI systems. Each system owner must ensure their system operates within its reviewed parameters, that monitoring is conducted as required, that incidents and anomalies are reported promptly, and

that material changes are assessed before implementation. System owners are named in the AI system inventory.

Lending officers hold human oversight responsibility for the pre-adjudication lending model. Every referral and every decline produced by the model is reviewed by a lending officer before a final decision is communicated to the member. Lending officers also handle member requests for human review.

All staff who interact with AI systems are responsible for using them within the bounds of this policy and the AI Acceptable Use Policy, and for reporting concerns or anomalies to the relevant system owner.

6. AI System Inventory

Meridian maintains a register of all AI systems in scope. The CRO is responsible for the accuracy and completeness of this register.

Each entry in the register must record at minimum:

- The system name and a plain-language description of its purpose
- The name of the system owner
- The data inputs used, including the sensitivity classification of each input under the Data Classification and Handling Standard
- The population affected and the nature of the outputs produced
- The current risk tier
- The date of the most recent impact assessment and the date the next review is due
- The hosting environment and, where applicable, the supplier name
- Any conditions or limitations placed on the system's operation

The register is reviewed by the CRO annually. The register is available to the Board Risk Committee on request.

At the date this policy takes effect, the AI systems in scope include the pre-adjudication lending model used to score consumer loan applications up to \$15,000. Any new AI system, or any system that comes to meet the definition in Section 3, must be added to the register before it is deployed.

7. Risk Classification

Every AI system in the register must carry a risk tier. The tier is assigned through the impact assessment process described in Section 8 and is maintained in the AI system inventory.

Risk tiers are defined in the AI Risk Management Standard. At a minimum, the tier reflects:

- The likelihood that the system produces an incorrect or harmful output
- The severity of the impact on members if it does, including financial harm, denial of access to credit, and reputational harm
- The potential for disparate impact on identifiable member groups, assessed as a distinct impact category
- The reversibility of the harm
- The volume and vulnerability of the member population affected

The pre-adjudication lending model is classified as high risk by default, given that its outputs directly influence credit decisions for members. This classification may only be revised through a full impact assessment reviewed by the CRO and accepted by the Board Risk Committee.

A system's risk tier must be reassigned whenever a material change occurs or whenever monitoring evidence indicates that the system's actual risk profile differs from its assessed profile.

8. Impact Assessment and Pre-Deployment Review

No AI system may be deployed, and no material change to an existing system may be implemented, without a completed impact assessment.

The impact assessment must address:

- The purpose of the system and the business need it serves
- The data inputs, their sources, and their sensitivity classification
- The member population affected and the nature of the outputs
- The potential for harm, including disparate impact on member groups
- The controls that will govern the system, including human oversight arrangements

- The supplier arrangements, where applicable, and their alignment with the Third-Party Risk Management Policy
- The risk tier assignment and the rationale for it
- Any residual risks and the conditions under which they are accepted

The system owner prepares the impact assessment. The CRO reviews it. Where the assessment assigns a high risk tier, or where residual risk remains high after controls are applied, the Board Risk Committee must review and accept the risk before deployment proceeds. The review record is retained in accordance with the Records Retention and Disposal Policy.

For the pre-adjudication lending model, the impact assessment must specifically address the \$5,000 threshold below which approved applications proceed without a further human step. The assessment must document the basis for that threshold, the controls that compensate for the absence of human review at that tier, and the monitoring arrangements that will detect if the threshold produces disproportionate harm.

9. Human Oversight and Decision Rights

Meridian assigns human oversight to every AI system in proportion to its risk tier. The oversight arrangements for each system are documented in the impact assessment and recorded in the AI system inventory.

For the pre-adjudication lending model, the following oversight arrangements apply:

- Every application scored as a referral is reviewed by a lending officer before a decision is communicated to the member.
- Every application scored as a decline is reviewed by a lending officer before a decision is communicated to the member.
- Applications scored as approved where the loan amount is \$5,000 or below may proceed without a further human review step, subject to the controls and monitoring arrangements documented in the impact assessment for that threshold.
- Applications scored as approved where the loan amount exceeds \$5,000 require a lending officer review before the decision is communicated to the member.

Any member may request human review of any AI-influenced decision affecting them. That request must be granted. The lending officer conducting the review must have the authority to

reach a different outcome from the model. The outcome of the human review, and the basis for it, must be recorded.

The model may not be used for any purpose outside consumer loan application scoring up to \$15,000. It must not be used for collections decisions, account closure decisions, or any decision about an existing member relationship outside the loan application context.

Where a system owner or lending officer believes a system is producing outputs that are incorrect, biased, or otherwise harmful, they must escalate to the CRO immediately. The CRO may suspend a system's operation pending investigation.

10. Data Governance for AI Systems

AI systems at Meridian may only use data that is identified in the system's impact assessment and recorded in the AI system inventory. The data inputs for the pre-adjudication lending model are: application data, income verification information, internal account history, and credit bureau data used under the existing bureau agreement.

No member data may be transmitted to a third-party model provider for any purpose other than the scored transaction itself. No member data may be used by any vendor to train, fine-tune, or improve a model. These prohibitions must be reflected in every supplier contract and are monitored as part of the supplier oversight arrangements described in Section 12.

All member data used by AI systems is classified as restricted. Handling obligations are set out in the Data Classification and Handling Standard. Privacy obligations are set out in the Privacy and Personal Information Policy.

Member identity and contact data, income and employment information, account and transaction history, and credit bureau reports are all in scope. Processing of this data by AI systems must remain within Canada. The pre-adjudication scoring service operates in AWS ca-central-1. Core banking operates on premises in Ontario. No member data may be processed in any other region without a further impact assessment and Board Risk Committee review.

11. Ongoing Monitoring

Every AI system in operation must be monitored throughout its operational life. The system owner is responsible for monitoring and for acting on what monitoring reveals.

Monitoring must cover:

- **Performance:** whether the system continues to produce outputs that are accurate and consistent with its validated behaviour
- **Fairness and disparate impact:** whether the system's outputs show patterns of differential treatment across identifiable member groups, assessed annually and recorded in the monitoring plan for each system
- **Drift:** whether changes in the input data distribution or the member population are causing the system's outputs to diverge from its validated behaviour
- **Incidents and complaints:** whether members or staff have raised concerns about AI-influenced decisions, and whether those concerns reveal a systemic issue
- **Supplier changes:** whether the supplier has made changes to the model, its infrastructure, or its subprocessors that affect Meridian's risk assessment

Monitoring findings are reported to the CRO. Where monitoring reveals a material control failure, an emerging fairness concern, or a supplier change that affects the system's risk profile, the system owner must escalate to the CRO without delay. The CRO determines whether the Board Risk Committee must be informed.

The metrics used to assess AI governance performance across Meridian are:

- The proportion of AI systems in the register that carry a current impact assessment
- The proportion of high-risk AI uses that have a named human oversight arrangement in place
- The proportion of required periodic reviews completed within the scheduled timeframe
- The number of reviewed exceptions that have passed their expiry date without renewal
- The number of material control failures recorded and the number of associated corrective actions that are overdue

The CRO reports these metrics to the Board Risk Committee as part of the annual AI governance report.

12. Supplier Oversight

AI systems delivered or supported by third parties are subject to the Third-Party Risk Management Policy and the Third-Party AI Review Procedure. This section states the additional obligations that apply specifically to AI suppliers.

The scoring model vendor and the credit bureau are classified as critical suppliers. Any engagement with a critical AI supplier, and any material change to that engagement, requires Board Risk Committee review before it proceeds.

Every contract with an AI supplier must include:

- The right for Meridian to audit the supplier's AI-related controls
- Evidence of model validation, provided at a frequency set by Meridian and recorded in the contract
- A prohibition on using member data for training or model improvement
- An obligation to notify Meridian of a suspected breach or material model change within 24 hours
- Subprocessor approval requirements, such that the supplier may not engage a subprocessor that touches member data without Meridian's prior written agreement
- Business continuity commitments aligned with the Business Continuity Policy
- Defined exit assistance obligations to support transition or termination

Supplier performance against these obligations is reviewed by the system owner annually and recorded in the supplier register. Material concerns are escalated to the CRO.

13. Exceptions

A request to operate an AI system, or to use an AI system in a way, that does not comply with this policy must be submitted to the CRO as a formal exception request. The request must state the nature of the non-compliance, the business reason, the compensating controls proposed, and the period for which the exception is sought.

The CRO reviews the request. Where the exception relates to a high-risk system or involves a residual risk the CRO considers material, the request is referred to the Board Risk Committee.

Every reviewed exception is time-limited. The expiry date is recorded in the AI system inventory. The system owner is responsible for ensuring the exception is renewed or the non-compliance is remediated before the expiry date. Exceptions that pass their expiry date without renewal are reported to the CRO and counted in the governance metrics in Section 11.

14. Retirement and Decommissioning

When an AI system is retired, the system owner must ensure:

- The system is decommissioned in a manner that prevents unauthorised access to member data it held or processed
- Member data is disposed of or retained in accordance with the Records Retention and Disposal Policy
- The AI system inventory is updated to reflect the retirement date and the method of decommissioning
- Any supplier contract is terminated in accordance with its exit assistance provisions
- Evidence of the impact assessments, monitoring records, and review records produced during the system's operational life is retained in accordance with the Records Retention and Disposal Policy

The CRO confirms that retirement has been completed and that the inventory and records are in order.

15. Compliance and Enforcement

Compliance with this policy is mandatory for all staff and contractors within its scope. Breaches of this policy are managed through Meridian's standard disciplinary and performance processes.

Supplier non-compliance is managed through the contractual mechanisms described in Section 12 and the Third-Party Risk Management Policy.

The CRO monitors compliance and reports material failures to the Board Risk Committee. Where a failure poses an immediate risk to members, the CRO may suspend the affected AI system's operation pending remediation.

16. Relationship to Other Documents

This policy sits at the top of Meridian's AI governance framework. The documents below it implement its requirements:

- **AI Risk Management Standard** — sets the mandatory, testable requirements for risk tier assignment, impact assessment content, monitoring thresholds, and control specifications
- **AI Acceptable Use Policy** — governs how individual staff may use AI tools
- **Third-Party AI Review Procedure** — sets out the steps for reviewing and onboarding AI suppliers
- **AI Incident Response Procedure** — sets out the steps for detecting, containing, and reporting AI-related incidents
- **Privacy and Personal Information Policy** — governs the handling of member personal information, including when processed by AI systems
- **Third-Party Risk Management Policy** — governs supplier risk management, including AI suppliers
- **Records Retention and Disposal Policy** — governs how long records produced under this policy must be kept
- **Data Classification and Handling Standard** — defines the sensitivity classifications applied to data used by AI systems
- **Business Continuity Policy** — sets the continuity requirements that AI supplier contracts must reflect
- **Information Security Policy** and **Access Control Policy** — govern the security controls that protect AI systems and the data they use

Where this policy and another document address the same topic, the more specific document governs the operational detail. Where there is a conflict, the CRO determines which instrument takes precedence and records that determination.

17. Review and Maintenance

The CRO reviews this policy annually and whenever a material change to Meridian's AI use, its risk environment, or applicable legal or regulatory requirements makes an earlier review necessary. The review is recorded.

The Board Risk Committee reviews and accepts this policy following each review. The review record identifies the version, the date, and the names of those who participated in the review.

All staff within scope are notified when this policy is updated.

18. Effective Date and Version

This policy takes effect on the date it is reviewed and accepted by the Board Risk Committee. The version number, effective date, and next scheduled review date are recorded on the document cover sheet maintained by the CRO.

Owner review items

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

Type	Detail	Where
gap	The \$5,000 threshold below which approved loan applications proceed without a human review step is a significant governance decision. The policy requires this threshold to be documented and justified in the impact assessment for the pre-adjudication lending model, but the impact assessment itself has not been supplied. The customer must ensure that assessment exists, addresses the threshold explicitly, and has been reviewed by the CRO and accepted by the Board Risk Committee before the model operates.	Section 8 — Impact Assessment and Pre-Deployment Review; Section 9 — Human Oversight and Decision Rights
gap	The organisation's jurisdiction is Ontario, Canada. Applicable privacy and financial services legislation may include Ontario's privacy legislation, federal privacy legislation, and sector-specific regulation applicable to credit unions in Ontario. The customer's legal reviewers should confirm which instruments apply and ensure this policy and the Privacy and Personal Information Policy are	General — jurisdiction and regulatory alignment

	consistent with those obligations. No statutory requirements have been stated in this policy because the supplied material does not contain them.	
gap	The monitoring cycle for fairness and disparate impact assessment is set to annual in this policy. The customer must confirm this cycle is recorded in the monitoring plan for the pre-adjudication lending model and reflected in the AI Risk Management Standard.	Section 11 — Ongoing Monitoring
gap	The supplier contract review cycle is set to annual in this policy. The customer must confirm this cycle is recorded in the supplier register and in relevant supplier contracts.	Section 12 — Supplier Oversight
gap	The AI system inventory review cycle is set to annual in this policy. The customer must confirm this cycle is recorded in the inventory itself.	Section 6 — AI System Inventory
gap	The policy refers to the AI Risk Management Standard for the detailed definition of risk tiers and the testable requirements for impact assessments. The customer must ensure that standard is consistent with the risk classification approach described in this policy, particularly the treatment of fairness and disparate impact as a distinct risk category.	Section 7 — Risk Classification; Section 8 — Impact Assessment and Pre-Deployment Review
vague	The organisation has not specified physical office locations. The policy states that core banking operates on premises in Ontario and that no member data may be processed outside Canada. If Meridian has staff or operations in locations other than Ontario, the data residency commitment in Section 10 should be reviewed to confirm it remains accurate.	Section 10 — Data Governance for AI Systems
gap	The Records Retention and Disposal Policy governs how long records produced under this policy must be kept. This policy does not state retention periods because the supplied material does not contain them and they are owned by that sibling document. The customer must confirm that the Records Retention and Disposal Policy covers AI governance records, including impact assessments, monitoring records, review records, and exception records.	Section 14 — Retirement and Decommissioning; Section 8 — Impact Assessment and Pre-Deployment Review

Information the organization must supply

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

- A completed impact assessment for the pre-adjudication lending model, addressing the \$5,000 auto-approval threshold, the controls that compensate for the absence of human review at that tier, and the monitoring arrangements. This assessment must be reviewed by the CRO and accepted by the Board Risk Committee before the model operates under this policy.

- Confirmation that the monitoring plan for the pre-adjudication lending model records the annual fairness and disparate impact assessment cycle and that this cycle is reflected in the AI Risk Management Standard.
- Confirmation that the annual supplier contract review cycle is recorded in the supplier register and in relevant supplier contracts.
- Confirmation that the annual AI system inventory review cycle is recorded in the inventory itself.
- Confirmation from legal counsel of which privacy, financial services, and credit union regulatory instruments apply to Meridian in Ontario, and confirmation that this policy and the Privacy and Personal Information Policy are consistent with those obligations.
- Confirmation that the Records Retention and Disposal Policy covers AI governance records, including impact assessments, monitoring records, review records, and exception records, and that the retention periods in that policy are appropriate for the regulatory environment.
- Confirmation that the AI Risk Management Standard reflects the risk classification approach in this policy, including the treatment of fairness and disparate impact as a distinct risk category, before this policy takes effect.
- The effective date and version number, to be recorded on the document cover sheet once the Board Risk Committee has reviewed and accepted this policy.
- Confirmation of Meridian's physical office locations, to ensure the data residency commitment in Section 10 is accurate.

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.