VINCENT VANYNH CONSULTING

HUMAN-CENTRED AI TRANSFORMATION

Designing Human Control for AI-Assisted Work

Five applied use cases, model test results, and practical guidance for responsible automation

Human control is not about preserving manual work. It is about defining where AI may act, where human judgment matters, how exceptions are handled, and who remains accountable.

Applied Use Case Report | June 2026

Vincent Vanynh, MBA, PMP, CMC

Executive Summary

A practical test of how organizations can increase automation without surrendering judgment, visibility, or accountability.

Many organizations are no longer asking whether employees will use AI. They are asking where AI should be embedded into real work, how much authority it should have, and what must remain under human control.

The difficult part is not choosing between "human" and "AI." The difficult part is designing an operating arrangement that applies human attention where it adds value while allowing AI to remove unnecessary effort, increase throughput, and improve access to information.

> **Central finding**
>
> **Human control is strongest when it is designed into the workflow through clear authority, validation, monitoring, escalation, and accountability. Requiring a person to click approve on every output is not automatically safer, and may create little more than approval theatre.**

This report tests the Human-Control Decision Model against five contrasting, hypothetical but realistic use cases:

1. AI-assisted executive and strategic analysis
2. AI-generated client-facing communication
3. An employee policy and knowledge assistant
4. High-volume administrative transaction processing
5. AI-supported employment screening and recommendation

Across the five tests, the model produced directionally sound results. It also exposed three areas requiring refinement: an eligibility gate for autonomous execution, a specific control for external organizational communications, and a distinction between nominal review and meaningful review.

Final model conclusion

The four Human Control Patterns remain valid: Human-Led, Human-Reviewed, Human-Supervised, and AI-Executed with Guardrails and Escalation. The scoring range should remain a starting point, but it must be constrained by eligibility rules, mandatory overrides, and an operational-readiness test.

Why Human Control Is Not Anti-AI

The objective is to use AI more confidently and more extensively where the work can support it.

Human control is sometimes interpreted as a demand to keep a person in every step of a process. That interpretation is too narrow. A well-designed AI-supported workflow may automate a significant amount of work while still remaining under effective human control.

The question is not simply whether a human is present. The question is whether the organization has intentionally defined:

- The outcome AI is intended to improve
- The actions AI is permitted to take
- The boundaries it may not cross
- The evidence or criteria used to validate quality
- The conditions requiring escalation or intervention
- The person or role accountable for the outcome
- The ability to pause, override, or redesign the arrangement

> **A better objective**
>
> **Use AI to reduce avoidable effort, expand human capability, and perform suitable work at scale. Preserve human attention for judgment, exceptions, contested decisions, relationships, and consequences that require accountable discretion.**

This approach is consistent with established responsible-AI guidance. The NIST AI Risk Management Framework is designed to help organizations incorporate trustworthiness into the design, use, and evaluation of AI systems. Canada's Directive on Automated Decision-Making recognizes that automation may be partial, including situations where a system provides an assessment, recommendation, score, or summary to a human decision-maker. OECD principles likewise call for human agency and oversight mechanisms appropriate to context.

The Human-Control Decision Model does not replace those frameworks. It translates their broad expectations into a narrower operating question: how should responsibility be divided within a specific task, decision, or workflow?

Human control is - and is not

Human control is	Human control is not
Clear authority, boundaries, monitoring, escalation, and accountability.	A requirement that a person manually approves every output.
A way to direct human attention toward judgment and exceptions.	A reason to preserve low-value manual work.
Compatible with high levels of bounded, monitored automation.	A claim that AI should only retrieve information or draft text.

The Human-Control Decision Model

A workflow-level method for selecting the appropriate relationship between AI execution and human responsibility.

The model evaluates a defined work unit rather than an entire department, platform, or AI strategy. A work unit may be a decision, task, action, transaction, communication, or workflow segment.

Six oversight dimensions are scored from 1 to 4:

Consequence of error	**The impact of an incorrect output or action.**
Uncertainty	The ambiguity, variation, and reliability of available inputs and rules.
Reversibility	How easily an incorrect outcome can be corrected.
Failure detectability	How likely the organization is to notice an incorrect or inappropriate result.
Human judgment requirement	The degree to which context, discretion, professional expertise, ethics, or relationships matter.
Sensitivity and obligation	The extent to which the work involves confidential, regulated, personal, or rights-affecting information and outcomes.

The initial total produces a provisional pattern:

Score	Provisional pattern	Operating intention
6-9	AI-Executed with Guardrails and Escalation	AI completes defined low-risk work within approved boundaries.
10-14	Human-Supervised	AI performs recurring work; people monitor, sample, and handle exceptions.
15-19	Human-Reviewed	AI proposes an output or action; a person validates it before consequence.
20-24	Human-Led	AI supports preparation and analysis; a person independently decides and acts.

The provisional result is then tested against mandatory overrides and operational readiness. This prevents low scores in several dimensions from cancelling out one unacceptable risk or obligation.

Test Method

Each use case was assessed as an operating workflow, not as a generic technology concept.

For each case, the analysis followed the same sequence:

1. Define a specific work unit and business outcome.
2. Describe the proposed AI contribution.
3. Score the six oversight dimensions.
4. Identify the provisional Human Control Pattern.
5. Apply eligibility gates and mandatory override rules.
6. Test whether the organization can operate the pattern meaningfully.
7. Define permitted AI authority, required human responsibility, and minimum controls.

The examples are illustrative. They are intended to demonstrate the method, not to provide legal, privacy, cybersecurity, regulatory, employment, or technical advice for a specific implementation.

> **What makes the test useful**
>
> **The use cases were selected to create tension between value and control: ambiguous analysis, external communication, authoritative knowledge guidance, scalable routine processing, and consequential individual decisions.**

USE CASE 1

AI-Assisted Executive and Strategic Analysis

Using AI to accelerate research, synthesize evidence, identify patterns, and challenge assumptions without delegating executive judgment.

Scenario

An executive team is evaluating whether to consolidate service operations across several business units. The team wants an enterprise AI tool to summarize internal reports, compare options, identify risks, produce initial scenarios, and draft material for an executive decision session.

Business outcome	**Reduce preparation time and improve the breadth of evidence considered before an executive decision.**
Work unit	Prepare an initial strategic options analysis for executive review.
AI contribution	Retrieve, summarize, compare, structure, and challenge; produce a non-final options paper.
Downstream use	Executives assess trade-offs, consult stakeholders, and decide whether and how to proceed.
Primary risk	A persuasive but incomplete synthesis may overstate evidence, omit context, or anchor the executive team prematurely.

Oversight assessment

Dimension	Score	Level	Rationale
Consequence of error	3	Significant	A flawed analysis could influence investment, operating-model, workforce, or service decisions.
Uncertainty	4	Very high	The problem is strategic, contested, incomplete, and dependent on assumptions.
Reversibility	2	Reversible	The analysis can be corrected before a final decision, although early framing may create momentum.
Failure detectability	3	Difficult	A polished synthesis can appear credible even when evidence is incomplete or misweighted.
Human judgment requirement	4	Inherently human	Executives must balance competing objectives, relationships, risk appetite, and organizational context.
Sensitivity and obligation	3	High	The analysis may contain confidential strategy, financial, workforce, or stakeholder information.
Total	**19**	**Provisional range**	**Human-Reviewed**

Decision

Provisional result	**Human-Reviewed based on a score of 19.**
Override or eligibility finding	Human judgment is Level 4, triggering the inherently human judgment override.
Final Human Control Pattern	Human-Led
Readiness decision	Ready with conditions
AI may	Summarize approved sources, compare scenarios, identify contradictions, draft non-final analysis, and surface assumptions.
Human must	Independently evaluate evidence, add context, engage affected leaders, decide the recommendation, and own the decision.

Minimum controls

- Use approved and attributable information sources.
- Require traceability from major claims to source material.
- Separate evidence, assumptions, inference, and recommendation.
- Assign a human author or accountable analyst to challenge omissions and framing.
- Use AI to generate counterarguments and alternative scenarios, not only a preferred answer.
- Do not present AI-generated recommendations as objective or independently authoritative.
- Retain the executive team's normal decision process, consultation, and accountability.

What leaders should notice

A Human-Led pattern does not diminish the value of AI. AI can still remove a large share of information-gathering and first-draft effort. The value lies in expanding the decision-makers' field of view while preserving accountable judgment.

USE CASE 2

AI-Generated Client-Facing Communication

Using AI to produce project updates and routine correspondence without allowing speed to create inaccurate commitments or reputational exposure.

Scenario

A professional-services firm wants AI to draft weekly client status updates from project plans, issue logs, meeting notes, and delivery trackers. The objective is to reduce administrative effort and improve consistency across engagements.

Business outcome	**Reduce preparation effort while producing timely, consistent, and useful client updates.**
Work unit	Draft and issue a weekly client status communication.
AI contribution	Synthesize project records and produce a proposed external communication.
Downstream use	The client relies on the communication to understand progress, risks, decisions, and commitments.
Primary risk	The message may introduce an inaccurate status, implied commitment, inappropriate tone, or disclosure.

Oversight assessment

Dimension	Score	Level	Rationale
Consequence of error	3	Significant	Incorrect status or commitments may affect trust, commercial expectations, or project decisions.
Uncertainty	2	Moderate	The source records are structured, but context and interpretation vary.
Reversibility	2	Reversible	A message can be corrected, but trust and reliance effects may remain.
Failure detectability	2	Routinely detectable	A knowledgeable project lead can compare the draft with current records.
Human judgment requirement	2	Limited judgment	Most content is structured, but tone, emphasis, and commitments require judgment.
Sensitivity and obligation	3	High	The communication may include confidential, commercial, contractual, or relationship-sensitive information.
Total	**14**	**Provisional range**	**Human-Supervised**

Decision

Provisional result	**Human-Supervised based on a score of 14.**
Override or eligibility finding	External organizational representation requires a stronger pattern than the total alone suggests during initial deployment.
Final Human Control Pattern	Human-Reviewed
Readiness decision	Ready with conditions
AI may	Generate a draft from approved project records, highlight unresolved data, and propose standard wording.
Human must	Verify accuracy, remove inappropriate disclosures, confirm commitments, adjust tone, approve, and remain accountable for the communication.

Minimum controls

- Restrict source inputs to approved project records.
- Prevent the model from inventing dates, commitments, owners, or decisions.
- Visibly flag missing or inconsistent source data.
- Require approval by the project manager or accountable relationship owner.
- Use standard content sections and an approved tone guide.
- Retain a record of the approved final version and responsible sender.
- Evaluate whether selected low-risk sections can move to supervised automation after sustained performance is demonstrated.

Model refinement identified

The raw score placed this use case in Human-Supervised. That is too permissive for an initial client-facing workflow. The model needs an External Representation rule: communications that speak for the organization require human review unless they are tightly bounded, low-risk, and supported by demonstrated controls.

USE CASE 3

Employee Policy and Knowledge Assistant

Providing faster access to approved information while preventing a helpful interface from becoming an unaccountable source of policy interpretation.

Scenario

An organization wants an internal assistant that answers employee questions about leave, travel, procurement, benefits, and routine procedures. The assistant uses an approved knowledge base and is expected to reduce repetitive inquiries to support teams.

Business outcome	**Improve access to routine guidance and reduce time spent searching policies or waiting for support.**
Work unit	Respond to an employee's routine policy or process question.
AI contribution	Retrieve approved content, summarize it in plain language, and direct the employee to the relevant source.
Downstream use	The employee may act based on the answer or contact a specialist for clarification.
Primary risk	A confident answer may simplify, combine, or misstate a policy and lead an employee to act incorrectly.

Oversight assessment

Dimension	Score	Level	Rationale
Consequence of error	2	Limited	Most incorrect answers create rework or delay, though some topics may have greater impact.
Uncertainty	2	Moderate	Approved sources exist, but employee circumstances and policy exceptions vary.
Reversibility	2	Reversible	Most errors can be corrected through support processes.
Failure detectability	3	Difficult	Employees may not know the answer is wrong if it sounds plausible.
Human judgment requirement	2	Limited judgment	Routine retrieval is bounded; exceptions and interpretation require specialists.
Sensitivity and obligation	2	Moderate	The tool may receive employee context and provide guidance affecting internal actions.
Total	**13**	**Provisional range**	**Human-Supervised**

Decision

Provisional result	**Human-Supervised based on a score of 13.**
Override or eligibility finding	No mandatory override, provided the assistant remains retrieval-grounded and does not determine individual entitlement or resolve exceptions.
Final Human Control Pattern	Human-Supervised
Readiness decision	Ready with conditions
AI may	Answer routine questions from approved sources, cite the source, state limitations, and route exceptions.
Human must	Own policy content, monitor accuracy, review samples and escalations, correct the knowledge base, and handle interpretation or individual decisions.

Minimum controls

- Use only current, approved, version-controlled source material.
- Display the source document and effective date with each substantive answer.
- Define topics the assistant must not interpret or decide.
- Escalate conflicting sources, exceptions, personal circumstances, and consequential questions.
- Monitor unanswered questions, low-confidence responses, corrections, and repeated misunderstandings.
- Provide a simple way for employees to challenge or report an answer.
- Reassess topics individually; a general travel-policy answer and a benefit-eligibility determination should not use the same pattern.

> **What leaders should notice**
>
> **This is a strong candidate for supervised automation because the work is frequent, the knowledge can be bounded, and exceptions can be routed. The control model should be applied at the topic level so that routine guidance is not confused with consequential interpretation.**

USE CASE 4

High-Volume Administrative Transaction Processing

Allowing AI to complete routine work within clear tolerances while people supervise exceptions and operating performance.

Scenario

A finance team processes a high volume of standard invoices. The organization wants AI to extract invoice data, match it to purchase orders and receipts, apply routine coding, and prepare or post transactions that fall within defined tolerances.

Business outcome	**Reduce repetitive data entry and cycle time while directing staff attention to exceptions.**
Work unit	Process a standard, matched invoice within approved financial and data tolerances.
AI contribution	Extract, validate, match, code, and execute or prepare the transaction.
Downstream use	The transaction enters the financial workflow and may proceed toward payment.
Primary risk	Incorrect extraction or matching could create duplicate, inaccurate, or unauthorized processing.

Oversight assessment

Dimension	Score	Level	Rationale
Consequence of error	2	Limited	A bounded error creates financial rework or limited loss within established tolerances.
Uncertainty	1	Low	The standard case is defined by structured documents, matching rules, and tolerances.
Reversibility	2	Reversible	Transactions can be stopped, corrected, or reversed through established procedures.
Failure detectability	1	Immediately detectable	Validation rules, duplicate checks, matching exceptions, and reconciliations can expose failure.
Human judgment requirement	1	Rule-based	The eligible standard case follows explicit criteria.
Sensitivity and obligation	2	Moderate	Financial and supplier information requires protection and traceability.
Total	**9**	**Provisional range**	**AI-Executed**

Decision

Provisional result	**AI-Executed based on a score of 9.**
Override or eligibility finding	Passes the autonomous-execution eligibility gate because no dimension exceeds Level 2 and the standard case can be bounded.
Final Human Control Pattern	AI-Executed with Guardrails and Escalation
Readiness decision	Ready with conditions
AI may	Complete standard transactions within documented value, match, supplier, coding, and data-quality tolerances.
Human must	Approve boundaries, monitor performance, review exceptions and samples, investigate incidents, and retain authority to suspend the workflow.

Minimum controls

- Define the exact population of transactions eligible for autonomous processing.
- Set value, variance, supplier, duplicate, tax, account, and confidence tolerances.
- Route all out-of-boundary cases to a qualified person.
- Maintain segregation of duties and approved financial authority.
- Retain transaction evidence, system actions, exceptions, and interventions.
- Reconcile aggregate results and review statistically meaningful samples.
- Use a rapid pause or rollback process when error thresholds are exceeded.

What leaders should notice

This is the pattern that demonstrates the method is not anti-automation. Human control is maintained through bounded authority, monitoring, exception management, auditability, and intervention - not by requiring someone to approve every standard invoice.

USE CASE 5

AI-Supported Employment Screening and Recommendation

Using AI to organize evidence without allowing a score, ranking, or recommendation to become the de facto employment decision.

Scenario

A large organization receives hundreds of applications for recurring roles. It is considering AI to compare applications with job criteria, summarize candidate evidence, identify missing information, and recommend which candidates should proceed to interview.

Business outcome	**Reduce administrative review time and help recruiters examine a larger candidate pool consistently.**
Work unit	Assess application evidence and determine whether a candidate should advance.
AI contribution	Extract qualifications, compare evidence with criteria, summarize, score, rank, or recommend.
Downstream use	The output can determine whether an individual receives further consideration for employment.
Primary risk	A plausible ranking may embed inappropriate assumptions, omit transferable evidence, or effectively deny an opportunity without meaningful review.

Oversight assessment

Dimension	Score	Level	Rationale
Consequence of error	4	Severe	An incorrect result can deny employment consideration and create legal, fairness, and reputational consequences.
Uncertainty	3	High	Applications vary widely and may not express comparable experience in standardized terms.
Reversibility	3	Difficult to reverse	A missed candidate may never be reconsidered and the opportunity may pass.
Failure detectability	3	Difficult	Bias, proxy effects, or missing context may not be obvious in a polished score or summary.
Human judgment requirement	4	Inherently human	Assessment requires contextual interpretation, fairness, role knowledge, and accountable discretion.
Sensitivity and obligation	4	Critical	The workflow uses personal information and affects an important individual interest.
Total	**21**	**Provisional range**	**Human-Led**

Decision

Provisional result	**Human-Led based on a score of 21.**
Override or eligibility finding	Severe consequence, inherently human judgment, individual rights, sensitivity, and difficulty of reversal all reinforce the strongest pattern.
Final Human Control Pattern	Human-Led
Readiness decision	Not ready for AI ranking or autonomous recommendation; limited support use may proceed with controls
AI may	Extract and organize job-related evidence, identify application sections, prepare neutral summaries, and flag information for human examination.

Provisional result	**Human-Led based on a score of 21.**
Human must	Define criteria, assess candidates, consider context and transferable evidence, make advancement decisions, document reasoning, and provide appropriate challenge or correction paths.

Minimum controls

- Do not allow an AI score or rank to automatically remove candidates from consideration.
- Limit AI to evidence organization unless legal, privacy, fairness, and technical specialists validate a broader use.
- Require recruiters or hiring managers to review the underlying candidate evidence, not only the AI summary.
- Test for systematic differences, omissions, and proxy effects before and during use.
- Provide a process for correcting inaccurate information or challenging an outcome where appropriate.
- Document the business owner, decision authority, review method, and permitted AI contribution.
- Reassess the workflow whenever criteria, role requirements, candidate populations, or models change.

What leaders should notice

Adding a human after an AI ranking does not automatically create meaningful oversight. If the ranking determines attention, ordering, or exclusion, it may already be shaping the employment decision. The workflow must preserve the human decision-maker's ability and obligation to examine the underlying evidence.

Cross-Case Results

The model differentiated between support, review, supervision, and bounded execution in a way that reflects how the work actually operates.

Use case	Score	Provisional	Final pattern	Primary reason
Executive analysis	19	Human-Reviewed	Human-Led	Inherently human strategic judgment
Client communication	14	Human-Supervised	Human-Reviewed	External organizational representation
Policy assistant	13	Human-Supervised	Human-Supervised	Bounded knowledge with escalation
Invoice processing	9	AI-Executed	AI-Executed	Low uncertainty and bounded reversible action
Employment screening	21	Human-Led	Human-Led	Rights, sensitivity, judgment, and consequence

The total score was useful, but two cases demonstrated why it cannot be the sole decision mechanism. Executive analysis required Human-Led control because one dimension - inherently human judgment - was decisive. Client communication required Human-Reviewed control because the output represented the organization externally, even though its total fell within the supervised range.

What passed the test

- The four patterns were sufficiently distinct to classify all five workflows.
- The dimensions captured both operational and human consequences.
- The model permitted meaningful automation in the policy and transaction cases.
- Mandatory overrides prevented false precision from the total score.
- The readiness assessment exposed situations where a suitable pattern could not yet be operated credibly.

What needed refinement

- Autonomous execution requires a positive eligibility gate, not only a low total score.
- External communications need a specific representation rule.
- Review must be tested for competence, information, capacity, and authority.
- Workflows should be segmented so that routine cases and consequential exceptions do not inherit the same control pattern.
- A review date and change triggers should be mandatory because AI performance, source information, and operating context change over time.

Revised Selection Logic

The test supports retaining the scoring model while strengthening the conditions around it.

Refinement 1 - Autonomous execution eligibility gate

AI-Executed with Guardrails and Escalation should only be available when every oversight dimension is Level 1 or 2, no mandatory override applies, the eligible work population is explicitly bounded, and failure can be detected through reliable controls.

Refinement 2 - External representation rule

AI-generated content that represents the organization to a client, employee, public audience, regulator, partner, or other external party should begin as Human-Reviewed unless the content is tightly templated, low-risk, based on approved information, and has demonstrated sufficient performance to justify supervised operation.

Refinement 3 - Meaningful review test

A Human-Reviewed pattern is only valid when the reviewer has competence, access to evidence, sufficient capacity, and authority to reject or change the output. Where those conditions do not exist, the process must be redesigned rather than labelled human-reviewed.

Refinement 4 - Workflow segmentation

One business process may require more than one pattern. Routine cases may be AI-Executed, exceptions Human-Reviewed, and disputed or consequential cases Human-Led. The assessment should be performed at the smallest practical work unit.

Refinement 5 - Reassessment triggers

Every approved pattern should include a review date and event-based triggers, including changes to the model, data, policy, population, workflow, law, vendor, error rate, or consequence profile.

> **Revised rule of thumb**
>
> **Use the score to establish a provisional range. Use eligibility gates to permit automation. Use overrides to protect consequential interests. Use readiness tests to determine whether the organization can operate the design in practice.**

Ten Questions for an Executive or Project Team

A practical discussion guide for reviewing an AI-supported workflow before pilot or scale.

1. What exact task, decision, or action are we asking AI to support or perform?
2. What business outcome should improve, and how will we know?
3. What may AI do independently, and what is outside its authority?
4. What would happen if the output were wrong but appeared credible?
5. Which cases are routine, and which require judgment or exception handling?
6. Who can validate the output, and do they have the evidence, time, expertise, and authority to do it?
7. What conditions require escalation to a person?
8. Who can override, pause, or stop the workflow?
9. Who remains accountable for the final outcome?
10. What evidence would justify moving to more automation - or moving back to stronger review?

Suggested use

Choose one live AI opportunity and answer these questions with the business owner, users, technology team, and relevant governance specialists. Where the answers differ, the organization has identified a design issue before it becomes an operational problem.

From Discussion to Practical Implementation

The model is designed to create an actionable workflow decision, not a general statement of AI principles.

Organizations can apply the model at several levels:

Single use case	**Define the permitted AI contribution, Human Control Pattern, controls, and pilot conditions for one task or workflow.**
Functional workflow	Segment routine, exception, and consequential work across multiple control patterns.
Pilot portfolio	Compare proposed use cases and identify which are ready, which need stronger controls, and which should not proceed.
Executive governance	Establish common language for authority, human responsibility, escalation, and accountability across initiatives.
Implementation recovery	Diagnose an AI workflow that is creating rework, unclear ownership, weak adoption, or uncontrolled risk.

A focused Human Oversight and AI Workflow Review would typically produce:

- A defined work-unit and business-outcome statement
- An oversight assessment and final Human Control Pattern
- A current-to-future workflow map
- Permitted AI authority and human responsibilities
- Validation, monitoring, escalation, and intervention requirements
- Role, capability, and adoption impacts
- Pilot conditions and a prioritized implementation action plan

> **Practical next step**
>
> **Select one AI use case that is important enough to matter but contained enough to examine. Map how the work is performed today, define what AI would change, and decide the Human Control Pattern before selecting the final workflow or launching the pilot.**

Conclusion

Confident AI adoption requires more than access, policy, or a nominal human in the loop.

Organizations create value from AI when they are clear about where it fits into real work. That requires more than identifying a promising use case. It requires deliberate decisions about authority, judgment, workflow, validation, escalation, monitoring, capability, and accountability.

The five use cases show that human control and automation are not opposing objectives. The same model that requires Human-Led control for an employment decision can support AI-Executed processing for a bounded administrative transaction. The difference is not enthusiasm for AI. The difference is the nature of the work and the organization's ability to control it.

The practical goal is not to keep people performing every task. It is to make people faster and more capable while keeping outcomes, exceptions, and accountability under intentional human direction.

About the Author

Transformation, governance, execution, and human-centred AI adoption.

Vincent Vanynh, MBA, PMP, CMC, is a transformation consultant whose work spans enterprise implementation, governance, process and role design, stakeholder alignment, operational readiness, training, and adoption. His Human-Centred AI Transformation practice helps organizations move from scattered experimentation toward practical AI-supported workflows with clear ownership, appropriate oversight, and sustainable operating discipline.

This applied use case report forms part of the Human Oversight and AI Workflow Design Toolkit.

Contact and resources

vanynh.com | Tools and Resources

Important limitation

This report provides a business, workflow, governance, and adoption framework. It does not constitute legal, regulatory, privacy, cybersecurity, employment, safety, financial, or technical-model advice. Organizations should involve the appropriate specialists for their jurisdiction, sector, use case, and technology.

Selected references

National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). Source

Government of Canada. Guide on the Scope of the Directive on Automated Decision-Making. Source

Government of Canada. Algorithmic Impact Assessment tool. Source

OECD.AI. Human-centred values and fairness principle. Source

OECD.AI. AI Principles overview. Source