

FROM AI INTEREST TO ADOPTION CAPACITY

A practical agenda for people, organizations and communities

*Introducing the MANUAL Method: six moves, three scales and evidence gates for useful, responsible
AI adoption*

THE PROPOSITION You already know your people, your mandate and the work that matters. New Manual builds the AI adoption program around them.

About this paper

New Manual is a practical AI adoption, implementation and capacity-building practice based in British Columbia and working across Canada. It helps people, organizations and communities move beyond awareness and isolated experimentation toward useful, responsible ways of applying AI in real work.

This white paper sets out a research-grounded point of view and introduces the MANUAL Method, New Manual's applied framework for assessing adoption capacity, choosing worthwhile opportunities, piloting safely and building the capability to sustain what works.

METHOD NOTE This is an integrative practice paper, not a systematic review or a legal opinion. It draws on peer-reviewed experiments, implementation science, technology-adoption research, official Canadian statistics, labour and skills research, standards and public-sector guidance available as of 24 August 2026. The MANUAL Method is a structured diagnostic and implementation method; it is not a psychometrically validated scale, certification or guarantee of results.

Contents

Executive summary

1. The adoption gap
2. What the research says
3. Adoption capacity: the missing unit of progress
4. The MANUAL Method
5. The Adoption Capacity Map
6. Choosing, piloting and scaling use cases
7. A call to action for people
8. A call to action for organizations
9. A call to action for communities and networks
10. Starting pathways and partnership model

Appendices and references

Executive summary

AI adoption is no longer a future question. It is already occurring through formal programs, vendor rollouts, employee experimentation and everyday consumer tools. The harder question is whether that activity becomes dependable capability - and whether the resulting changes improve the work, the organization and the lives of the people affected.

THE CANADIAN SIGNAL In the second quarter of 2026, 19.2% of Canadian businesses reported using AI to produce goods or deliver services, triple the 6.1% reported two years earlier. In March 2026, 35.9% of workers reported using generative AI at work in the previous year. Individual use is moving faster than formal organizational adoption, while rural business use (9.9%) remains less than half the urban rate (21.0%). [1][2]

The evidence does not support a simple story in which access to a capable model reliably produces value. Controlled studies show large gains on some tasks: professional writing became faster and higher quality in one experiment; customer-support agents became more productive, with the largest gains among less experienced workers; and consultants improved on tasks that fell within AI's capability frontier. The same research also shows heterogeneity, misplaced reliance and worse performance when people use AI outside that frontier. [3][4][5]

The practical implication is decisive: adoption must be managed at the level of tasks, decisions and operating conditions. A licence, policy or workshop may be useful, but none is adoption on its own. Adoption exists when people can repeatedly use an approach in real work, under defined conditions, with acceptable risk, clear responsibility and evidence that it helps.

The MANUAL Method in one page

- Map the work and the outcome that matters.
- Assess adoption capacity across six conditions and three nested scales.
- Navigate value, risk and readiness; choose a route, including 'do not use AI.'
- Undertake a bounded pilot on real work with a baseline, review protocol and evidence contract.
- Anchor what proves useful in roles, routines, tools, safeguards and training.
- Learn from outcomes, monitor change and renew the system as work and technology evolve.

The assessment looks at six conditions: purpose and value; task and decision fit; people, agency and inclusion; data, tools and infrastructure; accountability and assurance; and learning, evidence and sustainment. It examines those conditions at three scales: the person or practice, the organization or system, and the community or ecosystem.

CALL TO ACTION Do not begin by asking which AI tool to buy. Begin by identifying one real outcome, mapping the work that produces it and assessing the conditions required to improve it. New Manual's role is to make that process rigorous, participatory and practical - then help build the internal capability to continue it.

1. The adoption gap

Fast use, shallow integration

The speed of diffusion can disguise the shallowness of integration. Statistics Canada's March 2026 Labour Force Survey supplement found that, among workers who had used generative AI, 63.5% used it for some but not most tasks and 24.9% used it for almost no tasks. Only a minority used it broadly. [2] This is consistent with international evidence showing high self-reported firm adoption but limited time spent by executives and little measured productivity impact to date. In a 2026 survey of almost 6,000 executives across four countries, roughly 70% said their firms actively used AI, yet more than 80% reported no impact on productivity or employment over the preceding three years. [6]

These figures are not contradictory. They measure different layers: awareness, individual use, firm-level use, workflow deployment and realized outcomes. Recent diffusion research explicitly separates those layers because adoption at one does not imply adoption at another. [7] A worker can use a general-purpose assistant while the organization lacks approved tools, data rules, quality controls or a way to learn from that use. An organization can purchase a platform without changing a workflow. A community can host workshops while leaving access, relevance and implementation support unresolved.

The theatre of adoption

Many adoption efforts stop at visible activity: keynote sessions, tool demonstrations, prompt libraries, licences, hackathons or high-level policies. These can create awareness and momentum, but they are weak proxies for changed practice. Implementation science warns that innovations can appear promising in controlled settings and still fail because context, recipients, organizational conditions and wider systems work against implementation. [8][9]

A PRACTICAL TEST If the people doing the work cannot explain when the AI-assisted workflow should be used, what information may enter it, how outputs are checked, who remains responsible and what evidence would justify continuing, the organization has activity - not yet adoption capacity.

Unevenness is a design constraint

Adoption is uneven by industry, occupation, firm size, geography, education, access and perceived relevance. In 2026, Canadian business use ranged from 42.3% in information and cultural industries to 4.5% in agriculture, forestry, fishing and hunting. Rural business use was 9.9%, compared with 21.0% in urban areas. Workers in professional, scientific and technical services used generative AI far more often than those in manual and place-based sectors. [1][2]

A credible national or community strategy therefore cannot be a single introductory workshop delivered to everyone. It must segment audiences by work, context, exposure, constraints and ambition; provide different on-ramps; and invest in trusted intermediaries that can translate general capability into locally relevant practice. The G7's 2025 SME AI Adoption Blueprint similarly emphasized trusted intermediaries, practical guidance, broadband, skills and advisory support. [10]

2. What the research says

New Manual's framework is built from ten recurring findings across AI field experiments, technology-adoption theory, implementation science, adult learning, risk management and technology diffusion.

1. The task is the correct starting point

Jobs, organizations and communities are bundles of activities, decisions and relationships. AI capability is 'jagged': it can perform some apparently difficult tasks well and fail at nearby tasks that look similar. In the BCG field experiment, consultants using AI were faster and produced higher-quality work inside the frontier, but were more likely to make errors on a task outside it. [5] Exposure therefore cannot substitute for task analysis.

2. Effects are heterogeneous

Average gains conceal who benefits, under what conditions and at what cost. Customer-support agents gained 14% productivity on average, with much larger gains for novice and lower-skilled workers. [3] In a professional-writing experiment, time fell 40% and quality rose 18%. [4] These results are important, but they do not generalize automatically across work types, tools or populations.

3. Technology is only one part of a socio-technical system

The Technology-Organization-Environment tradition, UTAUT, CFIR and NASSS all point to context: perceived usefulness and effort, social influence, enabling conditions, leadership, workflow

compatibility, recipients, organizational capacity and the wider system. [8][9][11][12] The practical lesson is to assess the system around a tool, not merely the tool.

4. Complementary assets determine whether potential becomes value

Earlier technology waves and current AI research show that returns depend on co-investment in data, process redesign, management, training and organizational capital. OECD research links adopter productivity advantages to digital infrastructure, ICT capability and management and problem-solving skills. [13][14] A model can be inexpensive while implementation remains costly.

5. Participation is an operating requirement

People are not passive recipients of implementation. Consultation surfaces tacit knowledge, real constraints, workarounds and risks that leaders or vendors will miss. It also affects legitimacy and willingness to use the resulting system. ILO evidence finds substantially greater willingness where AI is introduced with employee consultation and clear guidance. [15] Participation is both a design input and a trust mechanism.

6. Training works when it transfers into practice

AI literacy matters, but generic knowledge fades unless people have relevant tasks, protected practice time, feedback, peer support, access to appropriate tools and permission to use what they learned. Research on training transfer consistently identifies opportunity to perform, the work environment and supervisor or peer support as important. [16] Programs should be designed around application, not attendance.

7. Trust must be calibrated, not maximized

Users may under-trust or over-trust AI. Explanations alone do not reliably prevent overreliance; cognitive forcing - such as requiring an independent judgment before seeing an AI recommendation - can reduce it, though sometimes with usability trade-offs. [17] Knowledge workers also report less critical-thinking effort when confidence in AI is high. [18] Human review must be designed into the task.

8. Responsible adoption is lifecycle work

NIST's AI Risk Management Framework organizes action around govern, map, measure and manage. ISO/IEC 42001 treats AI as a management system requiring continual improvement. Canadian privacy regulators emphasize legal authority, necessity and proportionality, accountability, safeguards, transparency and meaningful human challenge mechanisms. [19][20][21] A one-time policy is insufficient.

9. Scaling is a distinct implementation problem

A good pilot is not a small deployment. Scaling introduces more users, variability, data, integrations, edge cases, dependencies and consequences. NASSS research finds that complexity across multiple domains predicts non-adoption, abandonment and difficulty sustaining technology-supported change. [9] Scale should be earned through evidence and readiness, not assumed from enthusiasm.

10. Community adoption depends on intermediaries and shared capacity

SMEs and community organizations often lack specialist talent, procurement leverage and time to evaluate vendors. OECD research identifies neutral intermediaries, peer learning, demonstrations and 'test before invest' support as key diffusion channels. [22] Community programs need trusted conveners, shared resources and local relevance, not only content delivery.

SYNTHESIS The durable unit of progress is adoption capacity: the ability to identify worthwhile applications, implement them under appropriate conditions, learn from evidence and change course when the technology is not helping.

3. Adoption capacity: the missing unit of progress

A working definition

Adoption capacity is the demonstrated ability of a person, organization or community to choose, use, govern and improve AI-enabled ways of working so that they produce worthwhile outcomes under acceptable conditions.

This definition is deliberately stricter than access, awareness, intention or use. It includes the ability to decline an application, contain risk, preserve human responsibility, adapt a workflow and sustain learning after an external advisor or program has ended.



Figure 1. Adoption capacity is nested. A weakness at an outer or inner scale can constrain the others.

Three nested scales

Scale	Primary question	Typical conditions	Success looks like
Person / practice	Can I use AI well in the work that matters to me?	Agency, skill, judgement, access, boundaries, feedback	Repeatable task improvement without loss of responsibility or capability
Organization / system	Can we turn useful experiments into governed operating practice?	Leadership, workflows, roles, data, tools, policy, review, measurement	Documented, supported workflows with accountable owners and measurable outcomes
Community / ecosystem	Can benefits diffuse across people and organizations fairly and sustainably?	Trust, access, intermediaries, shared infrastructure, local relevance, peer learning, legitimacy	A portfolio of locally useful applications and durable delivery capacity

Capacity is a profile, not a single score

New Manual does not reduce readiness to a universal percentage. A composite score can average away a decisive weakness: a high-value use case with no lawful data route, no review capability or unacceptable consequences is not made safe by strong leadership elsewhere. The assessment therefore produces a profile, highlights critical constraints and defines a target state proportionate to the intended use.

The target is also not 'maximum maturity.' A solopreneur using AI for low-risk drafting needs a different operating system from a ministry using AI to support decisions about services. The right target depends on consequence, scale, sensitivity, reversibility and the ability of affected people to understand or challenge what happens.

4. The MANUAL Method

The MANUAL Method is New Manual's six-move process for turning AI interest into adoption capacity. It integrates task-based AI evidence, contextual diagnosis, human-centred design, controlled experimentation, implementation science and continuous improvement.




Figure 2. The MANUAL Method. Progress is gated by evidence; stopping is a valid outcome.

Move	What happens	Minimum artifacts
M — Map	Map the mandate, outcome, work, information, decisions and people affected.	Outcome statement; task and decision map; baseline; stakeholder map
A — Assess	Assess adoption capacity across the six conditions and three nested scales.	Capacity profile; target state; constraints; critical gaps
N — Navigate	Compare use cases by value, risk and readiness; choose the appropriate route.	Prioritized portfolio; risk tier; pilot, prepare, park or stop decision
U — Undertake	Run a bounded pilot on controlled real work with explicit review and measurement.	Pilot charter; approved data route; test set; evidence contract; incident log
A — Anchor	Embed what proves useful in roles, routines, tools, safeguards, documentation and training.	Standard workflow; decision rights; job aids; training; ownership; support model
L — Learn	Measure outcomes, monitor drift, share learning and renew the system as conditions change.	Outcome report; adoption evidence; lessons; review cadence; scale, adapt or retire decision

Five operating rules

- No use case without an outcome. The desired change must be observable and connected to the mandate or the person's goals.
- No pilot without a baseline. If the current performance is unknown, improvement cannot be distinguished from novelty or expectation.
- No deployment without an owner. A human role remains accountable for the workflow, its controls, its review and its retirement.
- No scale without evidence. A demonstration proves possibility; a pilot tests a hypothesis; deployment requires operational evidence.
- No adoption theatre. Attendance, licences, prompts generated and hours of training are inputs, not final outcomes.

STOP RULE If AI does not improve the work, if the benefit is too small to justify the implementation burden, or if the risk cannot be reduced to an acceptable level, the method recommends not proceeding. A well-supported 'no' is a successful decision.

5. The Adoption Capacity Map

The assessment examines six conditions at each relevant scale. Evidence may come from interviews, observation, workflow artifacts, policies, data samples, tool configurations, surveys, performance records and structured workshops. Ratings are supported by evidence and accompanied by confidence and criticality, rather than treated as precise measurements.

Condition	Core questions	Evidence sought
1. Purpose and value	Is the intended outcome specific, worthwhile and aligned with the person's goals, the organization's mandate or the community's priorities?	Outcome statement; baseline; beneficiary; value hypothesis; non-AI alternative
2. Task and decision fit	Which steps involve generation, prediction, retrieval, transformation or judgement? Where is context tacit? What happens when the output is wrong?	Task map; decision rights; failure modes; review points; reversibility
3. People, agency and inclusion	Do people have voice, motivation, time, skill and support? Who gains, who bears burden and who could be excluded or deskilled?	Stakeholder input; capability profile; job-quality and accessibility impacts; training-transfer plan
4. Data, tools and infrastructure	Are information, access, security, interoperability, procurement and vendor conditions adequate for the use?	Data classification; approved environment; vendor assessment; integration and support needs
5. Accountability and assurance	Are legal authority, privacy, fairness, IP, safety, human oversight, challenge, incident response and records proportionate to impact?	Risk tier; control set; owner; review protocol; escalation and contestability route
6. Learning, evidence and sustainment	Can the group test, measure, share, adapt and retire the workflow as performance, needs or tools change?	Evidence contract; learning cadence; documentation; monitoring; sustainment resources

A five-level evidence rubric

Level	Meaning	Decision implication
0 — Unexamined	The condition is unknown, assumed or absent.	Do not infer readiness. Investigate or stop.
1 — Aware	The issue is recognized, but evidence, ownership or practice is limited.	Clarify, assign and prepare.
2 — Prepared	Minimum conditions and controls are designed and resourced for a bounded test.	Eligible for an appropriately scoped pilot.
3 — Working	The condition is demonstrated in real use with documented practice and usable evidence.	Eligible for controlled deployment if critical conditions are met.
4 — Adaptive	The condition is monitored, improved and sustained as the work and technology change.	Scale or transfer with continued monitoring.

How the assessment becomes a plan

1. Define the current profile and the evidence supporting it.
2. Define the target profile required by the chosen ambition and risk tier.
3. Identify critical gaps - especially any condition that can veto the use case.
4. Sequence actions by dependency: foundations before tools, controls before exposure, practice before scale.
5. Assign an owner, resource, measure and review date to each action.
6. Reassess after the pilot and at material changes in technology, workflow, data or policy.

6. Choosing, piloting and scaling use cases

The Value-Risk-Readiness decision

The goal of assessment is not to produce a heat map for its own sake. It is to make decisions. Every proposed use case is considered across three independent dimensions.

Value. Magnitude and likelihood of a meaningful improvement for the intended beneficiary, including quality, time, reach, cost, resilience or experience.

Risk. Potential severity and likelihood of harm, considering sensitivity, rights, safety, bias, security, IP, reputation, environmental burden, reversibility and contestability.

Readiness. Strength of the conditions required to test or operate the use, including people, task design, data, tools, controls, support and learning capacity.




Figure 3. Readiness and value determine the route; the risk gate can override any quadrant.

Risk routing

Route	Typical characteristics	Minimum response
Low consequence	Reversible; no sensitive data; limited audience; easy verification	Basic data boundary, competent review, source checks and record of the workflow
Moderate consequence	Business or service impact; confidential information; repeated use; some affected parties	Named owner, approved tool, structured evaluation, sampling, incident and escalation process
High consequence	Rights, benefits, employment, education, health, safety, finance or vulnerable groups may be affected	Formal legal/privacy/security/human-rights review; impact assessment; strong human decision authority; contestability; independent assurance where appropriate
Unacceptable or uncontrolled	/ No lawful or secure data route; no meaningful review; irreversible harm; opaque high-impact decision; benefits do not justify exposure	Do not pilot or deploy. Redesign the use or choose a non-AI approach.

The bounded pilot

A pilot is a controlled learning instrument, not a celebration of the technology. It should be large enough to expose real workflow conditions and small enough to contain failure. The pilot charter should specify:

- the problem, intended beneficiary and value hypothesis;

- the current workflow and baseline;
- the exact task boundary and what remains human-only;
- approved data, tools and users;
- known failure modes and required review;
- success, guardrail and stop criteria;
- the comparison design, sample and duration;
- ownership, incident response and documentation; and
- the decision to be made at the end: stop, adapt, extend, deploy or scale.

The evidence contract

Before the pilot begins, participants agree on what evidence will count. This reduces retrospective storytelling and helps protect human outcomes from being displaced by a single efficiency metric.

Evidence family	Example measures	Why it matters
Outcome value	Cycle time, service level, revenue, cost, reach, backlog, responsiveness	Tests whether the use changes the outcome that justified it
Quality	Accuracy, completeness, error severity, rework, consistency, expert rating	Prevents speed from hiding degraded work
Adoption and implementation	Eligible users, active use, task coverage, fidelity, support demand, abandonment	Distinguishes a good demo from a usable operating practice
Human impact	Workload, autonomy, confidence, skill, job quality, accessibility, client experience	Shows who benefits, who bears burden and whether capability is being built or eroded
Risk and assurance	Privacy/security events, policy exceptions, bias checks, unsupported claims, escalations, challenge outcomes	Tests whether safeguards work in practice
Sustainability	Total cost, owner capacity, vendor dependence, documentation quality, monitoring effort, environmental considerations	Tests whether the result can be maintained and adapted

Scale gates

- Gate 1 - fit: the task and intended outcome are clear, and a non-AI option has been considered.
- Gate 2 - permission: the data route, tool, responsibilities and risk controls are approved for the pilot.
- Gate 3 - evidence: the pilot meets pre-agreed value, quality, human-impact and risk thresholds.
- Gate 4 - operability: documentation, support, monitoring, training and ownership are ready beyond the pilot team.
- Gate 5 - sustainment: the system continues to deliver value, remains within risk tolerance and can be changed or retired.

7. A call to action for people

For an individual, AI adoption is not the accumulation of prompts. It is the development of judgement: recognizing suitable tasks, framing work, protecting information, testing outputs, integrating results and knowing when to stop. AI literacy research defines capability broadly enough to include critical evaluation and effective collaboration, while UNESCO's competency work emphasizes human agency, ethics and progressive movement from understanding to application and creation. [23][24]

What to assess

- Goals: What outcome matters now - time, quality, confidence, employability, service, learning or business growth?
- Task portfolio: Which recurring tasks consume attention? Which involve sensitive information, consequential judgement or human relationships?
- Current capability: Can you explain what the system does, frame a task, evaluate an output, check sources and recognize uncertainty?
- Boundaries: What information must not be entered? Which decisions must remain yours? What external rules apply?
- Learning environment: Do you have a real task, a safe tool, feedback, practice time and a peer or coach?
- Evidence: Can you show that the new practice improves an outcome without weakening your own skill or responsibility?

A 30-day personal adoption sprint

1. Week 1 - Map: choose one recurring, low-consequence task and record the current time, quality problems and desired outcome.
2. Week 2 - Test: try two or three approaches on non-sensitive examples. Compare against your own work or a trusted standard.
3. Week 3 - Calibrate: create a review checklist, identify common failure modes and decide what must remain human-only.
4. Week 4 - Anchor: document the workflow, repeat it on real tasks, measure the result and decide whether to adopt, adapt or stop.

INDIVIDUAL CALL TO ACTION Start with one task you already understand well enough to judge. Build a repeatable practice, not a collection of tricks. Keep a small evidence log: task, approach, time, quality, error, lesson and next change.

What New Manual can provide

New Manual can offer public workshops, targeted cohort programs, personal task and capability assessments, guided practice, office hours and portfolio-building. For job seekers and people adapting to workforce change, the emphasis should be on transferable judgement and demonstrable workflows - not claims of expertise based on tool familiarity alone.

8. A call to action for organizations

For an organization, the central challenge is conversion: turning distributed experimentation into a coherent portfolio of useful, governed work practices. The opportunity is significant, but the implementation burden is real. In 2026, 44.4% of Canadian AI-using businesses reported changing training or staffing practices; only 32.0% reported AI-related training for existing employees, with much lower rates among the smallest firms. Cybersecurity or privacy concerns were the most commonly reported barrier. [1]

The first leadership questions

- Which organizational outcomes matter enough to justify change?
- Where are staff already using AI, formally or informally?
- Which workflows combine high friction, adequate volume, usable information and reviewable outputs?
- Where would error affect rights, safety, employment, reputation or vulnerable people?
- What information classes and approved environments apply?
- Who owns AI-enabled workflows after the pilot?
- How will staff and affected people participate in design and evaluation?
- What evidence would justify scaling - and what would trigger a stop?

An organization-wide starting sequence

Step	Action
1. Establish the mandate	Name an executive sponsor and operational lead. Define the outcomes, risk tolerance and decision authority.
2. Listen and map	Consult staff, map current work and identify existing use, workarounds, concerns and opportunity areas.
3. Set safe boundaries	Issue usable interim guidance on tools, data, review, disclosure and prohibited uses while deeper work proceeds.
4. Build the portfolio	Score candidate use cases by value, risk and readiness. Select a small, diverse set of pilots.
5. Pilot on real work	Use approved environments, representative tasks, human review, baselines, guardrails and incident tracking.
6. Operationalize	Document the new workflow, assign owners, train for transfer, integrate support and update relevant policies.

Step	Action
7. Govern and learn	Maintain a use-case inventory, review outcomes and incidents, monitor vendors and retire weak or outdated practices.

Responsible adoption in the Canadian context

Organizations operate under existing law and policy; generative AI is not a law-free zone. Applicable obligations vary by jurisdiction, sector, organization and activity. Canadian privacy regulators emphasize legal authority, necessity and proportionality, safeguards, transparency, accountability, limiting sensitive information, impact assessment and meaningful human challenge. [21] British Columbia's public-service policy requires protection of sensitive information, review of outputs and inclusive, respectful use. [25] Federal guidance similarly stresses responsibility for privacy, security, quality and human decisions. [26]

For mature or higher-impact programs, NIST AI RMF, ISO/IEC 42001, ISO/IEC 23894 and ISO/IEC 42005 can provide useful reference points for governance, risk management and impact assessment. [19][20] They do not replace context-specific legal, privacy, security, procurement, labour, accessibility, human-rights or professional advice.

ORGANIZATION CALL TO ACTION Within 30 days, establish a cross-functional adoption team, map current employee use, choose three real workflows for assessment and publish interim boundaries that staff can understand. Do not wait for a perfect enterprise strategy before making current use safer and more learnable.

What New Manual can provide

New Manual can lead the readiness assessment, workflow discovery, staff consultation, use-case portfolio, tool evaluation, risk routing, controlled pilots, documentation, role-based training and outcome measurement. The objective is to leave behind an operating capability: people who can continue to choose, govern and improve AI-enabled work.

9. A call to action for communities and networks

Community-scale adoption is not organization-scale adoption with a larger registration list. A community contains different mandates, resources, infrastructures, languages, risk profiles and starting points. It also contains existing trust relationships that determine who participates, whose knowledge is heard and whether support is credible.

The role of the trusted intermediary

Associations, economic-development organizations, educational institutions, workforce bodies, foundations, government agencies and community networks can reduce search costs, counter vendor distrust, pool expertise and create peer-learning effects. OECD and G7 work increasingly identifies trusted intermediaries as a core diffusion channel for SMEs, particularly for emerging technologies where local firms lack evaluation capacity. [10][22]

The intermediary should not be expected to become an AI vendor. Its comparative advantage is audience knowledge, convening power, legitimacy and continuity. New Manual can bring the adoption method, facilitation, implementation support and evaluation architecture around the people and organizations the intermediary already serves.

What to assess at community scale

Condition	Community-level question
Mandate and public value	What shared outcomes matter - productivity, service access, workforce resilience, local innovation, inclusion or sector capability?
Segments and needs	Which groups differ by work, exposure, access, language, digital confidence, geography and risk?
Access and infrastructure	Do participants have connectivity, devices, approved tools, data and practical support?
Trusted delivery capacity	Which organizations can recruit, convene, coach, host, refer and sustain the work?
Rights, sovereignty and legitimacy	Who has authority over data and decisions? How will Indigenous rights, community protocols, privacy, accessibility and meaningful participation shape the program?
Portfolio and learning system	How will use cases, outcomes, risks and lessons be compared across a diverse cohort without forcing standardization?

Design principles for a cohort or community program

- Segment before designing. Build different pathways for different work and starting conditions.
- Recruit through trusted channels. Participation is a relationship, not a media buy.
- Combine training with implementation support. Add diagnostics, office hours, workflow coaching and peer learning.
- Use real work and local examples. Generic demonstrations do not establish relevance.
- Provide safe access. Where possible, pair guidance with approved tools, sandboxes or test data.
- Measure reach and distribution. Track who participates, who completes, who applies, who benefits and who is missing.
- Create reusable local assets. Train facilitators, document workflows and establish referral and escalation networks.
- Report outcomes honestly. Separate awareness, confidence, application, implementation and sustained results.

Data sovereignty and community authority

Programs involving First Nations must be shaped by Nation-specific laws, governance, worldviews and protocols. The First Nations principles of OCAP® - ownership, control, access and possession - support information governance and the path to First Nations data sovereignty; they are not a generic checklist to be applied by outside organizations. [27] Engagement must begin with authority, relationships and consent, not with a proposed technology or data set.

More broadly, communities should be able to understand what is being proposed, influence priorities, question assumptions and decide that an application does not fit. Federal guiding principles for AI in government explicitly call for prioritizing individuals and communities, controlled test environments, inclusive public engagement, training and attention to digital divides. [28]

COMMUNITY CALL TO ACTION Convene a 90-minute adoption roundtable with the people who understand the audience, the work and the constraints. Do not ask for use cases first. Ask what outcomes matter, what is already happening, who is not being served and what shared capacity would make responsible experimentation possible.

What New Manual can provide

New Manual can co-design and deliver a place-based or sector-based adoption program that combines opportunity assessment, segmented learning, controlled pilots, implementation support, peer exchange, office hours, facilitator development and outcome reporting. The program can serve a single cohort or build an enduring local adoption function.

10. Starting pathways and partnership model

Three practical entry points

Entry point	Best for	First engagement	Decision produced
Personal assessment practice	Solopreneurs, workers, job seekers, graduates, career changers	Task inventory, capability profile and guided real-work test	30- or 60-day learning and workflow plan
Organization diagnostic adoption	Businesses, nonprofits, governments, institutions and associations	Leadership alignment, staff consultation, workflow mapping and capacity assessment	Prioritized use-case portfolio and 90-day pilot roadmap
Partner-led program adoption	Economic-development, workforce, education, sector and community networks	Audience segmentation, ecosystem mapping and program co-design	Cohort or community program with delivery and evaluation model

A 90-day organizational pathway

Period	Focus	Core outputs
Days 1-30	Map and assess	Mandate; current-use scan; staff input; workflow inventory; capacity profile; interim boundaries
Days 31-60	Navigate and prepare	Prioritized portfolio; risk routes; pilot charters; approved tools/data; baselines; evidence contracts
Days 61-90	Undertake and decide	Controlled pilots; review and incident logs; outcome readouts; deploy, adapt, extend or stop decisions

A 6- to 12-month community pathway

1. Phase 1 - Listen and segment: establish outcomes, audience groups, partners, barriers and baseline conditions.
2. Phase 2 - Build foundations: facilitator capacity, access routes, safe-use guidance, referral networks and evaluation design.
3. Phase 3 - Learn through work: targeted sessions built around real tasks and sector or community realities.
4. Phase 4 - Implement: guided pilots, office hours, peer circles and workflow documentation.
5. Phase 5 - Share and sustain: outcome report, reusable assets, train-the-trainer, partner handoff and next-cohort design.

The partnership proposition

NEW MANUAL'S ROLE You already know your people, your mandate and the work that matters. New Manual builds the AI adoption program around them.

This partnership model protects what each party does best. The partner brings trust, audience knowledge, mandate, recruitment channels and local continuity. New Manual brings the assessment method, adoption research, facilitation, workflow design, controlled experimentation, capacity building and measurement. Participants bring the practical knowledge of the work and retain meaningful agency over what changes.

What success should be reported

- Reach: who participated, who did not and whether priority groups were represented.
- Capability: what participants can now understand, do, evaluate and explain.
- Application: how many people tested AI on a relevant task under appropriate conditions.
- Implementation: how many workflows were documented, owned and used beyond the session.

- Outcomes: what changed in quality, time, service, confidence, access, workload or other intended measures.
- Responsibility: what safeguards, human-review practices and escalation routes became operational.
- Sustainment: what local people, assets, communities of practice and review mechanisms remain.

Conclusion: build the capacity to choose

The most important capability in a period of rapid technological change is not the ability to adopt everything. It is the ability to choose well: to identify genuine value, understand the work, include the people affected, create appropriate conditions, test honestly, govern proportionately and learn faster than the technology changes.

That capability can be built. It begins with assessment, but it does not end with a score. It produces a plan, a portfolio, a set of responsibilities and a learning system. For a person, it becomes a working practice. For an organization, it becomes an operating capability. For a community, it becomes shared infrastructure for inclusive and responsible adaptation.

THE INVITATION Choose one meaningful outcome. Map the real work behind it. Assess the conditions. Then take the smallest responsible step that can produce useful evidence. That is how AI interest becomes adoption capacity.

Appendices

Appendix A. Diagnostic question bank

The following questions illustrate the evidence sought in a New Manual assessment. They are adapted to the scope, sector, risk and scale of each engagement.

Condition	Scale	Illustrative questions
Purpose and value	Person	What outcome do you want to improve, and how will you know? What non-AI approach could achieve it?
Purpose and value	Organization	Which strategic or service outcome does the use support? Who benefits? What baseline exists?
Purpose and value	Community	Whose priorities define value? How are benefits expected to be distributed?
Task and decision fit	Person	Which steps are repetitive or generative, and which require your lived experience or judgement?
Task and decision fit	Organization	Where are handoffs, exceptions, tacit knowledge and consequential decisions?
Task and decision fit	Community	Which tasks recur across participants, and which require local or sector-specific adaptation?
People, agency and inclusion	Person	Do you have skill, confidence, time, feedback and a safe chance to practise?
People, agency and inclusion	Organization	How are staff and affected people involved? What may change in workload, autonomy, roles or skill?
People, agency and inclusion	Community	Who may be excluded by access, language, disability, geography, cost, trust or program design?
Data, tools and infrastructure	Person	What information will you use, and is the tool appropriate for it?
Data, tools and infrastructure	Organization	Are data classification, security, procurement, integration, retention and vendor terms understood?
Data, tools and infrastructure	Community	Can participants access connectivity, devices, safe tools, technical support and relevant examples?
Accountability assurance	Person	What must you verify? What decision remains yours? What could harm someone if wrong?
Accountability assurance	Organization	Who owns the use case, approves risk, reviews outputs, handles incidents and responds to challenges?
Accountability assurance	Community	What laws, rights, sovereignty, protocols and public expectations shape the program?
Learning and sustainment	Person	How will you record lessons and protect the skills you still need to exercise?
Learning and sustainment	Organization	What will be measured, reviewed, updated and retired? Who has capacity to do it?
Learning and sustainment	Community	How will learning travel across the network, and what local capability will remain?

Appendix B. Minimum use-case card

- Name and owner
- Intended beneficiary and outcome
- Current workflow and baseline
- Task boundary and human-only steps
- Inputs, information classification and approved environment
- Expected value and non-AI alternative

- Failure modes, affected people and reversibility
- Risk tier, controls, review and escalation
- Pilot scope, comparison, duration and users
- Success, guardrail and stop measures
- Deployment requirements and sustainment owner
- Decision date: stop, adapt, extend, deploy, scale or retire

Appendix C. Evidence design cautions

- Do not generalize a task experiment to an occupation, organization or sector without additional evidence.
- Distinguish causal experiments from observational surveys, self-report and vendor case studies.
- Record model, version, settings, data conditions and date; capability and risk can change materially.
- Use representative real work, including difficult and boundary cases - not only polished demonstrations.
- Measure absolute quality and serious errors, not just averages or user preference.
- Look for distributional effects: novice/expert, language, disability, gender, geography, job type and other relevant groups.
- Consider learning and deskilling over time, not only immediate output.
- Treat absence of reported incidents cautiously when reporting and monitoring systems are weak.
- Separate implementation outcomes - acceptability, adoption, feasibility, fidelity, cost, reach and sustainment - from the service or business outcome itself. [29][30]

Appendix D. Research method and limitations

This paper used a rapid integrative review designed for practice synthesis. Searches prioritized: (1) peer-reviewed field and controlled experiments on generative AI and work; (2) foundational and updated frameworks for individual adoption, organizational adoption and real-world implementation; (3) systematic reviews and authoritative research on skills, labour, SMEs, diffusion and human-AI reliance; (4) official Canadian statistics, policy and privacy guidance; and (5) recognized risk and management standards.

The review is broad rather than exhaustive. AI evidence is unusually time-sensitive: tools, user familiarity, organizational practice and policy continue to change. Some of the newest 2025-2026 evidence is survey-based or in working-paper form. Results from short, bounded tasks do not establish

long-term organizational productivity, job quality or community outcomes. The references therefore support design principles and hypotheses to test; they do not eliminate the need for local evidence.

References

- [1] Statistics Canada. Analysis on artificial intelligence use by businesses in Canada, second quarter of 2026 (2026). [Source](#)
- [2] Statistics Canada. Use of generative artificial intelligence tools among Canadian workers, March 2026 (2026). [Source](#)
- [3] Brynjolfsson, E., Li, D., & Raymond, L. R. Generative AI at Work. Quarterly Journal of Economics 140(2), 889-942 (2025). [Source](#)
- [4] Noy, S., & Zhang, W. Experimental evidence on the productivity effects of generative artificial intelligence. Science 381, 187-192 (2023). [Source](#)
- [5] Dell'Acqua, F. et al. Navigating the Jagged Technological Frontier. Organization Science, Articles in Advance (2026). [Source](#)
- [6] Yotzov, I. et al. Firm Data on AI. Stanford Institute for Economic Policy Research working paper (2026). [Source](#)
- [7] Bonney, K. et al. The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks. NBER Working Paper 35141 (2026). [Source](#)
- [8] Damschroder, L. J. et al. The updated Consolidated Framework for Implementation Research based on user feedback. Implementation Science 17, 75 (2022). [Source](#)
- [9] Greenhalgh, T. et al. Beyond Adoption: Nonadoption, Abandonment, Scale-Up, Spread, and Sustainability of Health and Care Technologies. Journal of Medical Internet Research 19(11), e367 (2017). [Source](#)
- [10] G7 Industry, Digital and Technology Ministers. SME AI Adoption Blueprint (2025). [Source](#)
- [11] Venkatesh, V. et al. User Acceptance of Information Technology: Toward a Unified View. MIS Quarterly 27(3), 425-478 (2003). [Source](#)
- [12] Tornatzky, L. G., & Fleischer, M. The Processes of Technological Innovation (1990). [Source](#)
- [13] OECD. Fostering an inclusive digital transformation as AI spreads among firms (2024). [Source](#)
- [14] Roth, F. Intangible capital and labour productivity growth at the macro, meso, and micro levels: a review of the literature and insights for public policies. Science and Public Policy 52(6), 893-909 (2025). [Source](#)
- [15] International Labour Organization. Generative AI at work: What it means for jobs in Europe and beyond (2025). [Source](#)
- [16] Govaerts, N., & Dochy, F. Disentangling the role of the supervisor in transfer of training. Educational Research Review 12, 77-93 (2014). [Source](#)
- [17] Bucinca, Z., Malaya, M. B., & Gajos, K. Z. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI. Proceedings of the ACM on Human-Computer Interaction 5(CSCW1) (2021). [Source](#)
- [18] Lee, H.-P. et al. The Impact of Generative AI on Critical Thinking. CHI 2025 (2025). [Source](#)
- [19] National Institute of Standards and Technology. AI Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1 (2024). [Source](#)
- [20] ISO. ISO/IEC 42001:2023 - Artificial intelligence management systems (2023). [Source](#)
- [21] Office of the Privacy Commissioner of Canada et al. Principles for responsible, trustworthy and privacy-protective generative AI technologies (2023). [Source](#)
- [22] OECD. SME Technology Adoption in the United Kingdom: Main findings from the consultation process (2026). [Source](#)
- [23] Long, D., & Magerko, B. What is AI Literacy? Competencies and Design Considerations. CHI 2020 (2020). [Source](#)

- [24] UNESCO. AI competency frameworks for students and teachers (2024). [Source](#)
- [25] Province of British Columbia. Policy on the use of generative AI (updated 2025). [Source](#)
- [26] Government of Canada. Guide on the use of generative artificial intelligence (updated 2026). [Source](#)
- [27] First Nations Information Governance Centre. The First Nations Principles of OCAP®. [Source](#)
- [28] Government of Canada. Guiding principles for the use of AI in government (2024). [Source](#)
- [29] Proctor, E. et al. Outcomes for Implementation Research. *Administration and Policy in Mental Health* 38, 65-76 (2011). [Source](#)
- [30] RE-AIM. Reach, Effectiveness, Adoption, Implementation, and Maintenance framework. [Source](#)
- [31] OECD. Generative AI and the SME Workforce: New Survey Evidence (2025). [Source](#)
- [32] OECD. The Adoption of Artificial Intelligence in Firms (2025). [Source](#)
- [33] OECD. AI and skills (2026). [Source](#)
- [34] International Labour Organization. Generative AI and Jobs: A Refined Global Index of Occupational Exposure (2025). [Source](#)
- [35] World Bank. World Development Report 2026: The Promise of Artificial Intelligence (2026). [Source](#)
- [36] Pereira, E., Graylin, A. W., & Brynjolfsson, E. The Enterprise AI Playbook: Lessons from 51 Successful Developments (2026). [Source](#)
- [37] Doshi, A. R., & Hauser, O. P. Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances* 10(28) (2024). [Source](#)
- [38] Government of Canada. AI Strategy for the Federal Public Service 2025-2027 (2025; overview updated 2026). [Source](#)

About New Manual

New Manual helps organizations, networks and individuals turn AI interest into practical capability and working systems - without the theatre, shortcuts or pressure to adopt technology for its own sake. Its work begins with objectives, responsibilities and existing workflows, then combines assessment, practical training, controlled pilots, implementation support, governance and measurement.

BEGIN THE PROCESS Start with an adoption capacity assessment: understand where you are, define where you need or want to go and build an evidence-based plan to get there.