


TEN SECTORS / REAL WORK / USEFUL CHANGE

What does good look like?

Practical AI by sector.

Choose a useful opportunity. Build it into your work.
Check what is actually improving.



Alongside People

Advice and hands-on implementation for Melbourne companies and government organisations.

OCTOBER 2026

Find your starting point.

01 / Government	More capacity for public service: less time rebuilding information
02 / Accounting	Create capacity for client advice: reduce the chase before review
03 / Software	Ship work you can stand behind
04 / Automotive	Keep the next step clear
05 / Insurance	Improve claims service: give people a clearer file sooner
06 / Creative	Make room for the good idea
07 / Healthcare	More time for care: improve the administration around it
08 / Education	Give education staff more time for the work that needs them
09 / Manufacturing	Less chasing, clearer production handovers
10 / Logistics	Resolve delivery exceptions with less chasing

Reading this collection

Choose the paper closest to your work. Each includes practical opportunities, illustrative examples, responsibilities and a way to examine whether a change helps. Relevant professional and organisational decisions stay with the people authorised to make them.

The structured HTML reading editions are available at alongsidepeople.com/resources. This downloadable edition has selectable text, live links and bookmarks; use HTML for structured screen-reader navigation.

GOVERNMENT

More capacity for public service: less time rebuilding information

Alongside People / Advice and hands-on implementation

A practical AI workflow paper for Melbourne government teams.

Executive takeaway

Your organisation may already have AI experiments underway without a shared view of what good looks like. We help connect those efforts to public-service priorities and choose where to go next. The example below makes that conversation concrete; you do not need a project selected before speaking with us.

A service team can spend valuable time rebuilding information that already exists: checking which notice changed, finding the authoritative date and turning several documents into a usable staff update. AI may help with that preparation. The useful outcome is more capacity for service work, with reliable information and clear responsibility.

Begin with one factual workflow, authorised sources and a person who owns the result. Compare total preparation and checking effort with the existing approach. Keep administrative decisions, eligibility assessments and consequential advice outside the pilot described here. A faster draft is useful only if the recipient can rely on it after the agreed checks.

Our proposed approach combines advice and implementation. We establish where the work is getting stuck, build a bounded change and help the organisation decide whether it deserves routine use. This paper describes a practical experiment, not a claim of government approval or a promise of savings.

Find the service bottleneck

Useful discovery questions concern the work itself. Does a coordinator repeatedly compare published updates? Do frontline staff search several places for the current instruction? Does a reviewer rewrite a summary because conditions or dates were omitted? Does every handover depend on one colleague remembering where the authoritative material lives?

These are possible bottlenecks to investigate, not findings about Melbourne agencies. Follow a recent example with the people who prepared and received it. Separate waiting for approval from active drafting. Ask what would make the recipient's next action easier. A shared source list or better update template might solve much of the problem without AI.

Choose a workflow that can be observed within a bounded trial. Avoid starting with decisions about an individual's rights, benefits or access to services. Even preparatory work needs scrutiny when it could influence a consequential decision.

A fictional running example

Fictional example, not a client or government case study: A Melbourne-based government service team prepares a weekly staff reference note from published service notices. A coordinator finds changed pages, collects dates and drafts an update. A communications owner checks it before distribution. Staff use the note to locate the authoritative notice; it does not replace that source.

The proposed tool produces a draft with links, source dates and unresolved questions. It cannot publish, contact the public, decide eligibility or change service rules. Information absent from a source remains an explicit gap. Internal approval of the tool, workflow and information is required before any trial. The example reports no measured benefit.

What would make this worth continuing?

The following is a decision aid for the fictional trial, not an observed result. Agree the limits with the people responsible before testing.

- **Accepted result:** A checked staff reference note with current source links, dates, conditions and visible unresolved questions.
- **Who judges it:** The authorised communications owner, with the staff who rely on the note.
- **Hold and return to the manual route when:** A condition is invented, an authoritative source cannot be verified or a draft crosses into eligibility or service decisions.
- **Evidence for another step:** Compare total preparation, source checking and recipient clarification for notes of similar scope.

Record the actual organisation, relevant policy and authorised policy owner in the experiment brief. A Victorian agency, Commonwealth body and local council need their own applicability decision; a paper about government does not establish one for them.

Use the [one-workflow worksheet](#) to record the agreed question, evidence and next decision.

Build dependability through five stages

Understand. Enter with the recurring reference-note job and its owner. Map source discovery, drafting, review, approval and distribution. Establish which notices are authoritative and what staff need from the note. The benefit hypothesis is less searching and fewer clarification requests. Exit only with a process map, baseline, source/permission boundaries, a named approver and a testable question. If no one can resolve conflicting notices, fix that responsibility before introducing AI.

Try. Enter after the sponsor agrees the permitted public sources, tool, checks and stop conditions. Test synthetic or approved public-information examples, including an expired notice, a changed condition and an unreachable page. The hypothesis is reduced draft-preparation effort without losing dates, conditions or links. Exit with reviewed observations, corrections and checking time, plus an owner decision. Stop if a condition is invented or a source cannot be verified; prepare the note manually.

Repeat. Enter when the trial justifies another test. Use consistent source dates, instructions and review criteria. Ask a second staff member to follow the process. The hypothesis is fewer omissions across

repeated updates. Exit with reproducible instructions, exception records, a tested manual route and comparison against the baseline. A change that works only with tidy examples or its original author is still a trial.

Operate. Enter only after routine use in this scope is authorised. Name an owner and backup, control access and make failed runs visible. In the fictional team, a changed source list or new tool feature goes through review before use expands. The hypothesis is more consistent staff updates with less chasing. Exit with operating guidance, approved scope, visible exceptions, exercised recovery and an owner review. Pause if approval, ownership or source integrity changes.

Improve. Enter when routine observations support a review. Compare preparation, checking, correction and staff usefulness with the original purpose. The hypothesis is retaining useful changes while avoiding unhelpful expansion. Exit each review with a documented retain, revise, reduce or retire decision and its next trigger. A clearer template with less AI may be the better outcome. Broader delegation requires a fresh approval and evidence decision.

This is Alongside People's original workflow framework, **Understand / Try / Repeat / Operate / Improve**. It is not a certification, validated assessment or compulsory ladder. Human approval can remain at every consequential step.

Identify the jurisdiction before applying a rule

The Victorian [generative AI administrative guideline](#) applies to its specified public service bodies, public entities and personnel. For organisations within its scope, it requires agency-approved tools to be used ahead of publicly available tools, limits information entered into unapproved tools to public information and retains personnel accountability. Information entered into an approved tool must stay within the protective marking the organisation permits for that tool. Organisation and sector policies may impose higher thresholds.

OVIC's [enterprise generative AI guidance](#) addresses privacy and information-security preparation and review for its Victorian public-sector audience. Purchasing an enterprise product does not complete that work.

The Commonwealth [responsible AI policy, version 2.0](#) applies to non-corporate Commonwealth entities with specified exceptions. It has mandatory requirements and is effective from 15 December 2025. It is not automatically a policy for Victorian agencies or local councils. A council must establish its own applicable requirements; this paper makes no blanket claim about council coverage.

Ask the organisation's authorised privacy, security, records, legal and procurement specialists to determine applicability and acceptance. Our framework helps prepare that conversation; it does not provide legal assurance, permission to use information or supplier-panel eligibility.

A proposed 30-day pilot

This is a planning example, not a service delivery promise. Begin the clock after the sponsor has the authority, people and approved environment needed. Approval may take longer than the experiment; no deadline overrides it.

Days 1-7: Map the reference-note workflow, select examples and record baseline observations. Agree source currency, review responsibilities and the exact scope. Produce an experiment brief and a decision on whether to proceed.

Days 8-14: Prepare a draft workflow using synthetic or authorised public sources. Run difficult examples alongside normal ones. Review every output, record defects and adjust the instructions. Hold any result with unresolved source questions.

Days 15-23: Repeat in shadow mode, so existing work remains authoritative. Ask another staff member to use the instructions and exercise the manual fallback. Collect the effort of preparer, reviewer and recipient.

Days 24-30: Review the comparison and remaining limits. Deliver the evidence, proposed operating responsibilities and an owner decision to continue testing, authorise a bounded next step, revise or stop. Routine use is a separate decision.

Measure capacity without hiding work

The primary measure could be total staff effort per accepted reference note. Record preparation, source verification, review, correction and distribution separately. Pair it with end-to-end turnaround, unanswered recipient questions and source accuracy. Track omitted conditions, stale information and failed-source handling as guardrails.

Preserve the original examples and criteria. Explain differences in workload or source complexity. Sparse observations reveal what happened in the sample; they cannot prove organisation-wide productivity. Released time may support other work, but it is not automatically a cash saving.

Record data flows, retention, access, audit needs and supplier involvement before implementation. Use approved storage for evidence, and avoid putting personal or sensitive material into a tool simply to test its limits. Everyone should know who can pause the workflow and how staff receive a corrected note.

How our team helps

Alongside People provides practical AI advice and hands-on implementation. Paul Volpato, our founder and CEO, brings the work's purpose, people and technical choices into one conversation. Our team can help map the process, design the pilot, implement an agreed draft-and-review workflow and prepare checks, training and handover.

Your organisation retains authority over information, approvals and service outcomes. We agree scope and operating responsibilities rather than imply indefinite support. The useful first conversation starts with a recurring job, its owner and what would make the result more dependable.

Official references

Accessed 3 October 2026. Recheck before consequential application.

- [Victorian administrative guideline](#), updated 12 February 2025: minimum requirements within its defined scope.

- [OVIC enterprise generative AI guidance](#), updated 26 June 2026: regulator guidance for the specified Victorian audience.
- [Commonwealth responsible AI policy v2.0](#), effective 15 December 2025: mandatory policy within its defined Commonwealth scope.

If you are already trying AI and are unsure what deserves further investment, [tell us what you have tried and what you want to improve](#). We can help find the starting point together.

Read the structured HTML edition: alongsidepeople.com/resources/ai-for-government

ACCOUNTING

Create capacity for client advice: reduce the chase before review

Alongside People / Advice and hands-on implementation

A practical AI workflow paper for Melbourne accounting practices.

Executive takeaway

Your practice may already be trying AI and still be unsure what good use looks like across the team. We help you assess those experiments, decide what matters and choose a useful next move. You do not need to arrive with a process already picked out.

An accounting team may have capable practitioners and good software yet lose time collecting the information needed for a proper review. Missing documents, inconsistent descriptions and repeated follow-ups can leave experienced people rebuilding the same picture before they can exercise judgement.

AI may help prepare that picture: organise permitted information, link it to sources and draft questions for review. The intended outcome is more capacity for client service and a shorter path to a complete review pack. The practitioner still owns the analysis, advice and professional acceptance.

Start with information preparation, measure the entire job and keep tax positions, professional judgements and lodgements under authorised human control. Our team combines advice and implementation to make one bounded change reviewable. We do not promise a multiple of productivity, guaranteed turnaround or a reduction in headcount.

Look at work waiting to become reviewable

Useful discovery begins with a recent job that waited. Which information was missing? Was the client asked twice because the original request was unclear? Did a reviewer discover that two documents covered different periods? Did a junior colleague spend time interpreting a note that should have been clarified with the client?

These are investigation questions, not claims about every practice. Follow the job from engagement and information collection to review. Separate client waiting time from staff effort. Examine who prepares, who checks and who resolves uncertainty. Standard document requests, clearer upload instructions or structured checklists may help before AI is involved.

Choose a process with repeatable inputs and an identifiable reviewer. A trial should help that person see the evidence more clearly. It should not produce confident conclusions from incomplete records or move difficult judgement into a tool that cannot own it.

A fictional running example

Fictional example, not an Alongside People client: A Melbourne accounting practice prepares a monthly review pack for a small business. Staff receive an agreed list of documents and explanations, mark what has arrived and assemble questions for the accountant. The accountant verifies the records and decides what further work or advice is needed.

The proposed tool prepares a source-linked document index and a draft list of unresolved information requests. It distinguishes “not supplied”, “unclear” and “requires practitioner judgement”. It cannot declare a transaction deductible, determine a tax position, post entries, lodge a return or send a client request without review.

Initial examples are synthetic. Any move to real records requires the firm's approval and appropriate client-information permissions. The fictional practice has no reported savings or observed results.

What would make this worth continuing?

The following is a decision aid for the fictional trial, not an observed result. Agree the limits with the people responsible before testing.

- **Accepted result:** A practitioner-accepted document index and list of unresolved questions.
- **Who judges it:** The reviewing accountant, with feedback from the preparer and client-request owner.
- **Hold and return to the manual route when:** A source is invented, a period mismatch is hidden or an unresolved matter becomes a tax conclusion.
- **Evidence for another step:** Compare total preparation, checking and follow-up effort for packs of similar scope; record omissions and repeat requests.

A useful fictional flag reads: “The supplied statement covers a different period; the practitioner needs to confirm which record is required.” “This expense is deductible” is an unsupported conclusion outside the proposed tool's job.

Use the [one-workflow worksheet](#) to record the agreed question, evidence and next decision.

Five stages from preparation to dependable use

Understand. Enter with the review-pack job, its practitioner owner and the people gathering documents. Map the request, collection, indexing, queries and review. Agree what “ready for review” means. The benefit hypothesis is fewer avoidable follow-ups and less reconstruction by the reviewer. Exit with a process map, document requirements, baseline, permission boundaries and one question to test. If readiness depends on unwritten knowledge, make that knowledge explicit first.

Try. Enter after agreeing permitted data, tool, reviewer, acceptance criteria and stop conditions. Use synthetic packs containing duplicate documents, mismatched periods and incomplete explanations. The hypothesis is reduced preparation effort while preserving traceability. Exit with observed defects, checking time and the practitioner's proceed/revise/stop decision. Stop if the tool invents a document, treats an unresolved matter as settled or suggests an unsupported tax conclusion. Return to the existing checklist.

Repeat. Enter when a useful trial merits further work. Establish consistent folder structure, period labels, source references and request wording. Have another preparer follow the instructions. The hypothesis is more complete packs with fewer repeated requests. Exit with reproducible instructions, documented exceptions, a manual route and a baseline comparison. The practitioner verifies that the change helps the review rather than simply moving work from preparation into correction.

Operate. Enter only after routine use in the bounded scope is approved. Identify the practitioner owner, operational backup and people allowed to change templates or access records. Retain evidence of what the tool prepared and what people corrected. The hypothesis is steadier turnaround and more usable staff capacity. Exit with operating guidance, visible failures, authorised access, change control and exercised recovery. Pause if permissions lapse or staff cannot recover a pack manually.

Improve. Enter with enough observations for a meaningful review. Compare accepted packs, total staff effort, repeat questions and reviewer experience. The hypothesis is learning which preparation changes genuinely help clients and staff. Exit each review with a retain, revise, expand or retire decision and a next trigger. A structured form may prove better than generated requests. Broader actions need separate authority and evidence; dependable human review can remain the endpoint.

Understand / Try / Repeat / Operate / Improve is our original per-workflow discussion framework. It is not a professional standard, validated readiness assessment or certification. A practice can operate one workflow dependably and still be exploring another.

Professional responsibility and client information

The Tax Practitioners Board's final [AI guidance statement, TPB\(GS\) 55/2026](#), was issued on 22 July 2026. It concerns registered tax and BAS agents providing tax agent services. It retains practitioner responsibility and calls for professional judgement and output review. Its confidentiality discussion addresses permission before client information is disclosed to a third party, including AI configurations that involve such disclosure.

APESB's [technical alert on ethical AI use](#) reminds members of the Australian Accounting Bodies that their ethical obligations under APES 110 continue when using AI. Membership, service type and other applicable standards need to be identified; do not treat tax-agent guidance as the complete rulebook for all accounting work.

OAIC's [commercial AI privacy guidance](#) supports product due diligence and privacy assessment where personal information is involved. Determine the firm's actual obligations with appropriate advisers. This paper does not provide tax advice or legal assurance.

A proposed 30-day pilot

This is a planning outline, not a delivery commitment. Start only once the firm has approved the scope, environment and information. Client permissions, professional review and access decisions may take longer.

Days 1-7: Follow recent review-pack preparation, establish baseline effort and agree the definition of readiness. Create a synthetic example set, including missing and contradictory material. Document the restricted actions and approval route.

Days 8-14: Build the index-and-question draft. Practitioners check it against the source records, record defects and inspect whether uncertainty remains visible. Improve templates before adding more automation.

Days 15-23: Repeat in shadow mode using authorised examples. The existing process remains authoritative. Test whether another staff member can prepare and recover a pack using the instructions. Review every outgoing request; no unattended client messaging.

Days 24-30: Compare accepted packs, review effort and remaining problems. Produce a decision record, operating proposal and handover material. Choose continued testing, an authorised bounded implementation, revision or stopping. The end of the month does not automatically establish readiness.

Baseline and guardrails

The primary outcome could be total staff effort to produce a practitioner-accepted review pack. Record preparation, checking, correction and client-request effort. Track elapsed turnaround separately, including time awaiting client responses. Compare similar periods and pack complexity rather than attributing every change to AI.

Guardrails include source traceability, incorrect missing-document flags, repeated requests, period mismatches and unauthorised disclosure. Include practitioner confidence and the client's experience of the questions. Preserve agreed criteria; do not count a pack as complete by silently excluding difficult matters.

Keep estimates separate from observations and record the limits of a small sample. Time released is capacity to allocate, not automatically profit or cash saved. Financial amounts should remain tied to verified records and the firm's established checking processes; fluent wording cannot validate a balance.

Record vendor access, retention, storage, subprocessors, permitted uses and change notifications. Masking a name does not by itself establish that records are anonymous or authorised for processing. Keep the person approving the use accountable and the manual process available.

How our team helps

Alongside People provides advice and hands-on implementation. Paul Volpato, our founder and CEO, connects operational purpose with practical systems work. Our team can map the preparation bottleneck, design the evidence plan, implement an agreed workflow and prepare instructions, checks and handover.

The firm's professionals retain judgement and acceptance. We agree scope and ongoing responsibilities rather than imply a tax-agent service, professional qualification or open-ended support. Bring one recurring job and a representative example. Together we can decide what is worth trying and what should stay with a person.

Official references

Accessed 3 October 2026; recheck before consequential use.

- [TPB\(GS\) 55/2026](#), issued 22 July 2026: final guidance for registered tax/BAS agents, not the earlier exposure draft.
- [APESB AI technical alert](#), 31 October 2025: professional-body guidance concerning members' continuing ethical obligations.
- [OAIC commercial AI guidance](#): regulator guidance, not a comprehensive statement of all applicable obligations.

If you are already trying AI and are unsure what deserves further investment, [tell us what you have tried and what you want to improve](#). We can help find the starting point together.

Read the structured HTML edition: alongsidepeople.com/resources/ai-for-accounting

SOFTWARE

Ship work you can stand behind

Alongside People / Advice and hands-on implementation

A practical path to better software delivery with people and AI

Your team may already be using coding assistants and agents, yet still be asking what good looks like beyond a faster first draft. We help you make sense of those experiments and decide what to improve next. You can bring that uncertainty; we can find the right starting point together.

For Melbourne software businesses and internal development teams, the useful question is whether AI helps a worthwhile change reach users with less friction and dependable quality. Writing a first draft of code is one part of that journey. Understanding the request, reviewing the result, releasing it and supporting it are part of the work too.

Our starting point is to start with one bounded delivery problem and measure the whole journey. A tool that generates more changes can increase the review queue. A tool that helps people resolve ambiguity, produce evidence and keep changes small may improve the flow. The trial needs to reveal which is happening in your team.

Our five stages are **Understand, Try, Repeat, Operate, Improve**. Each produces evidence before a larger commitment.

Find the constraint behind the busy team

Consider where a request actually waits. A product manager may know the desired outcome but leave edge cases unstated. An engineer may spend a morning finding the relevant behaviour in an unfamiliar service. A reviewer may need to reconstruct the purpose of a large change. A release may wait because nobody owns the operational check.

These are different problems. A coding assistant could help explore an unfamiliar repository. A clearer acceptance brief could remove repeated clarification. A standard evidence packet could make reviews easier. Some bottlenecks need a changed responsibility or a simpler process rather than an AI tool.

DORA's guidance examines throughput and instability across five delivery measures. It recommends application-level context and warns against targets people game. That supports measuring reliable delivery rather than code volume. [DORA delivery metrics](#).

Before adding anything, trace several recent changes from request to release. Separate active work from waiting. Ask the people doing the work what they repeatedly explain, check or repair. Their answers define the trial.

A fictional workflow: a clearer booking confirmation

This example is fictional and reports no client results.

Imagine a Melbourne business with a small booking product. Support receives questions because the confirmation page gives an unclear next step after a booking changes. The team chooses this one behaviour as a trial. It excludes payments, identity, data migration and any broad redesign.

The product owner writes a brief explaining the customer problem, the intended message, cancellation behaviour and the situations that must remain covered. An engineer gathers the relevant files, existing tests and design conventions. The AI assistant is asked to describe the current flow with file references before proposing a change. Where the explanation conflicts with the actual code, the engineer corrects it.

The assistant then prepares a small change and draft tests in an isolated working area. It must show what changed, why, and which checks actually ran. A successful test cannot become evidence for a different untested scenario. The engineer examines the diff, checks the expected behaviour against the brief and runs the relevant checks independently.

A reviewer receives the brief, the proposed change and the evidence together. They can reproduce the original issue and the corrected behaviour. If the change requires excessive explanation or touches unrelated code, it returns for simplification. The release owner then follows the normal release process and verifies the page after release.

GitHub's application card notes that generated code and explanations can be inaccurate and calls for review and testing. Its code-review documentation describes automated feedback and suggested fixes. We would use those capabilities as inputs to this workflow, while keeping a named person responsible for acceptance. [Copilot application card](#), [Copilot code review](#).

Shared evidence lets everyone assess the purpose and result. The trial would establish whether AI makes that journey easier.

What would make this worth continuing?

The following is a decision aid for the fictional trial, not an observed result. Agree the limits with the people responsible before testing.

- **Accepted result:** An accepted bounded change with reproducible evidence and a named release owner.
- **Who judges it:** The product owner, engineer, independent reviewer and release owner.
- **Hold and return to the manual route when:** A required check is weakened, a material regression is unresolved or the evidence claims a scenario that was not checked.
- **Evidence for another step:** Compare clarification, implementation, review and correction across similar changes; continue observing sparse production outcomes.

A useful fictional evidence packet names the requirement, describes the changed behaviour, lists checks actually run, identifies any remaining gap and names the release owner. Passing tests support the scenarios they exercise; they do not establish every possible product behaviour.

Use the [one-workflow worksheet](#) to record the agreed question, evidence and next decision.

Understand: agree what better means

We begin with the people responsible for the product, engineering, review and operation. Together we map one change type, identify its most costly friction and define an acceptable result. The baseline includes clarification, implementation, review, correction and release work.

The aim is a focused problem everyone recognises. Continue when there is a short workflow map, a written acceptance brief, a baseline sample and a named decision owner. We also record which repositories and information may enter the chosen tool. A tool decision follows that context.

Try: complete one change with evidence

We configure a contained trial around a familiar, reversible change. The team records the tool, model where visible, settings and information supplied. It compares the proposed result with the acceptance brief and keeps a record of failed attempts as well as useful ones.

The aim is learning from a complete task. Continue when there is a reviewable change, meaningful checks, reviewer notes and total human effort. A plausible explanation or a green screenshot alone is insufficient. If the team cannot reproduce the result, the trial is not ready to repeat.

Repeat: see whether the method survives variation

One well-chosen example can conceal a fragile process. We repeat the method across several comparable changes, including an ambiguous request and a case where the assistant should ask for clarification. We keep the scope small enough to inspect every result.

The aim is a reusable way of working. Continue when there is an agreed task brief, examples of accepted and rejected outputs, and evidence about review effort across the sample. A method that only works with one expert continually rescuing it needs more work before others depend on it.

Operate: make ownership and recovery ordinary

Once the method is useful, we connect it to the team's normal work. We establish access, task boundaries, review responsibility, cost visibility and a manual fallback. Release authority stays explicit. The team can stop a task, inspect its changes and return to an earlier state.

The aim is an improvement people can rely on during a normal week. Continue when there is a working runbook, an assigned owner, a demonstrated recovery and evidence from the approved environment. Passing a local trial does not establish that production access or monitoring is ready.

Improve: use results to choose the next change

We review whether the original constraint moved. Perhaps implementation became quicker while review became slower. Perhaps clearer briefs helped more than code generation. Either finding is useful. The next action might be narrowing the task, improving context, changing tools or stopping the approach.

The aim is informed investment. Continue when there is a dated decision supported by the trial record, with one next experiment and its success criteria. Expanding automation is an option when the evidence supports it.

A proposed 30-day trial

This proposed evaluation is not a commercial commitment or delivery guarantee.

During days 1-5, select one application and one recurring change type. Review recent comparable work, agree acceptance criteria and establish the baseline. During days 6-12, complete a contained example with an engineer and independent reviewer. Use days 13-21 to repeat the method, recording exceptions and correction work. During days 22-30, assess whether the method is suitable for routine use and document the decision.

If permissions, suitable examples or review capacity are unavailable, use synthetic material and describe the limits. Do not rush production integration to fit the calendar.

Measure completed value and its full cost

Agree thresholds before the trial. Compare similar changes and show the sample size. Record elapsed request-to-acceptance time, active human effort, clarification rounds, review time, material corrections and reopened work. Separate waiting for a release window from preparing the change.

Use quality guardrails alongside speed. Required checks must remain intact; known regressions and security findings must be resolved under the team's normal standards. Record any escaped defects and recovery work. For a short trial, production failure rates may be too sparse to support a conclusion. Continue observing rather than reporting a precise improvement from a tiny denominator.

Calculate cost per accepted change from tool usage, setup, human review and correction. Show recurring costs separately from the initial investment. Keep model or configuration changes visible so two runs are not treated as equivalent without explanation. Avoid individual rankings based on prompts, tokens or lines written.

The decision is straightforward: continue when the work becomes easier without weakening acceptance; revise when benefit depends on hidden effort; stop when the workflow adds more burden than it removes.

How we help

Alongside People brings advice and implementation together. We can help your team identify the constraint, define the task and its evidence, select an appropriate setup, and build the workflow around existing delivery practices. We then help evaluate what happened and make the operating method understandable to the people who will own it.

Paul Volpato is our founder and CEO. A useful starting conversation is about one recent change that took longer, required more review or created more rework than it should have.

Sources and scope

Understand / Try / Repeat / Operate / Improve is Alongside People's original discussion framework. It is not a certification, validated maturity assessment or official standard.

The workflow and five stages are our proposed approach. The example is fictional. Sources support the specific factual statements above and do not establish results for your business. All sources were checked on 3 October 2026; product features and access should be refreshed before a trial.

- [DORA's software delivery performance metrics](#), DORA; maintained guidance, checked 3 October 2026.
- [Application card: GitHub Copilot Chat](#), GitHub; maintained documentation, checked 3 October 2026.
- [About GitHub Copilot code review](#), GitHub; maintained documentation, checked 3 October 2026.

If you are already trying AI and are unsure what deserves further investment, [tell us what you have tried and what you want to improve](#). We can help find the starting point together.

Read the structured HTML edition: alongsidepeople.com/resources/ai-for-software

AUTOMOTIVE

Keep the next step clear

Alongside People / Advice and hands-on implementation

Better coordination for Melbourne dealerships and service businesses

You may already be trying AI across sales, service or administration and still be unsure what good looks like for the business. We help assess those experiments and choose a useful next move with your team. The service example below is one possibility, not a prerequisite for a conversation.

A customer should be able to understand what is happening with their vehicle, what needs their decision and when they will hear from you again. Your team should be able to answer those questions without reconstructing the job across several screens, notes and conversations.

Our starting point is to begin with that coordination work. AI may help prepare a status summary, organise an enquiry or draft a clear update. The useful result is a more dependable handover between customers, advisers, workshop staff and administration. Technical assessment and decisions about vehicle safety remain with the appropriately qualified people and established procedures.

Trial one operational journey through **Understand, Try, Repeat, Operate, Improve**. Establish whether the change clarifies the next step and reduces administration.

Follow the job between the departments

The service journey contains several different kinds of information. A customer describes a concern. A technician records findings. Parts staff confirm availability. An adviser explains the proposed work and asks for a decision. Accounts prepare the final record. Each person needs enough context to act without inheriting an unsupported assumption.

Friction can appear at any boundary. A promised callback has no owner. An arrival date is copied from an old message. An estimate is treated as a confirmed total. A customer's question reaches the workshop without the original context. A loan-vehicle arrangement lives in a conversation rather than the job record.

Sales operations have similar boundaries: an enquiry, an appointment, a vehicle's actual availability, a trade-in discussion and a delivery checklist. Different tasks need different controls. A drafting tool should not invent stock availability, finance suitability or a commercial promise.

Trace one journey before selecting a tool. An explicit status field and callback owner may solve the problem; an assistant may help assemble information.

A fictional workflow: the afternoon service update

This fictional example reports no client results.

Imagine a service department preparing afternoon updates for customers whose vehicles remain on site. The trial begins with synthetic job records. Each record separates technician findings, customer-authorised work, pending decisions, confirmed parts information and the next communication time.

An assistant prepares a short adviser packet for each job. Every factual line points to a dated source record. Missing information becomes a question. If parts availability has not been confirmed, the packet says confirmation is needed. If an additional repair is proposed, it remains pending rather than appearing in the completed-work summary.

The adviser checks the packet against the job record and technician's notes. They resolve discrepancies with the responsible person and prepare the customer update. The update explains what is known, what remains uncertain and the next step. It avoids technical conclusions that the technician has not supplied.

When a decision is needed, the adviser explains the proposed work and estimate through the business's authorised process. The decision and its scope are recorded in the system of record. A draft message never becomes evidence of customer authorisation by itself. The workshop acts on the authorised job record.

Consumer Affairs Victoria's guidance recommends written estimates describing work, parts, costs and timing, and advises customers to require contact before unauthorised repairs begin. That makes a clear distinction between information, estimates and permission especially useful in this workflow. [Car services and repairs](#).

The adviser sends the checked update and records the next callback owner. The packet can also be prepared manually. The trial measures preparation, checking and correction.

What would make this worth continuing?

The following is a decision aid for the fictional trial, not an observed result. Agree the limits with the people responsible before testing.

- **Accepted result:** A checked adviser packet separating confirmed status, proposed work and pending customer decisions.
- **Who judges it:** The service adviser, with the workshop and callback owner.
- **Hold and return to the manual route when:** An unconfirmed arrival, repair approval or price is presented as settled, or an unsupported promise reaches a customer.
- **Evidence for another step:** Compare preparation, checking and clarification for comparable service jobs; include callback commitments and duplicated contacts.

In a fictional packet, "Parts received: confirmed in the dated stock record" belongs under confirmed status. "Completion tomorrow: awaiting technician confirmation" belongs under pending questions. Neither the draft nor its wording authorises additional work.

Use the [one-workflow worksheet](#) to record the agreed question, evidence and next decision.

Understand: find the recurring missed handover

We map the selected journey with the people who perform it. Where does information enter? Which record is authoritative? Who may confirm availability, approve an estimate, speak to the customer and move the job forward? We examine recent examples and identify a recurring delay or ambiguity.

The aim is a shared picture of the work. Continue when there is a workflow map, named owners, clear status definitions and a baseline sample. We also identify information that should stay outside the trial. Customer details and workshop records need an approved handling arrangement before live use.

Try: prepare a checked packet for one job type

We start with a limited task, such as preparing an adviser update for an existing service booking. Synthetic examples include missing dates, conflicting notes and additional work awaiting a customer decision. The assistant must preserve those distinctions and surface questions.

The aim is a visible result advisers can judge. Continue when there is a completed packet, its source links, reviewer corrections and preparation time. A fluently written message with an unsupported promise is a failed example. We test the difficult cases deliberately.

Repeat: make the handover usable across shifts

We repeat the method with comparable records and different advisers. A shift change is a useful test: can another authorised person understand the job, identify the pending decision and continue the conversation? The method should not depend on the original author's memory.

The aim is a repeatable handover. Continue when there is a simple template, accepted and rejected examples, and a record of exceptions. Separate new delays from ones already present in the process. If the assistant repeatedly misreads a status, fix the input or narrow the task before adding volume.

Operate: connect carefully to everyday systems

After the trial shows value, we help define how the workflow fits the dealer or workshop management system. Begin with the least access needed. Make review and sending responsibilities explicit. Record who changes job status and who owns the customer's next update.

The aim is reliable coordination during normal operations. Continue when there is a tested procedure, approved access, cost visibility and a demonstrated manual fallback. Integration availability depends on the actual system and permissions; a demonstration does not establish a production connection.

The OAIC's guidance recommends due diligence when selecting AI products and highlights human oversight, information access and privacy risks. We use that as a starting point for checking the proposed information flow with the business's responsible advisers. It is not a certification of any particular tool or dealership. [OAIC commercial AI guidance](#).

Improve: fix what the evidence reveals

Review the exceptions alongside the routine work. Perhaps the packets save preparation effort but advisers still spend too long chasing parts information. Perhaps a shared status definition removes more friction than the assistant. Perhaps customers need fewer, clearer updates rather than more messages.

The aim is a focused next investment. Continue when there is a decision to continue, revise or stop, supported by the trial record. Extend the method to sales appointments or delivery preparation only after assessing their distinct sources, commitments and responsibilities.

A proposed 30-day trial

This proposed evaluation is not a commercial commitment or promised result.

During days 1-5, choose one department and one handover. Establish status definitions, responsibilities and the baseline. During days 6-12, test synthetic packets with advisers and include conflicting or incomplete records. During days 13-21, repeat the method; use limited live records only after the appropriate access and information-handling approvals. During days 22-30, review quality, staff effort and operating readiness, then document the decision.

Fit the trial to the team's checking capacity. If live use is not approved, use synthetic records and state what remains unproven.

Measure clarity, effort and promises kept

Before starting, define a comparable sample by job type and stage. Measure preparation and checking time, missing required fields, duplicated contacts, callback commitments completed within the agreed window and handovers returned for clarification. Record the count and denominator for every rate.

Track material errors separately: an incorrect status, unsupported completion date, wrong amount or confusion between proposed and authorised work. Those are more consequential than a minor wording correction. Pause the trial if such errors reach customers or the workflow changes authorisation without the required decision.

The baseline should include manual work and customer follow-up, not just message drafting. Compare routine and exception jobs separately. Parts delays, staffing changes and unusual workshop demand can affect the comparison. A short observation period may show a useful direction without establishing a causal business result.

Calculate cost per accepted packet from software use, preparation, checking and rework. Ask advisers whether the process makes the next action easier to identify. Ask customers a simple clarity question through an approved process if appropriate. Do not treat more messages or higher upsell value as proof of better service.

How we help

Alongside People combines advice with implementation. We can map the journey, improve the handover, assess a tool against the actual task and build a checked workflow around your existing systems. We help establish the baseline, review exceptions and prepare a method your people can operate.

Paul Volpato is our founder and CEO. Start with one recurring handover that makes your team chase information or leaves a customer unsure of the next step.

Sources and scope

Understand / Try / Repeat / Operate / Improve is Alongside People's original discussion framework. It is not a certification, validated maturity assessment or official standard.

The five stages and trial are our proposed approach. The example is fictional. This paper concerns dealership and service administration; it makes no diagnostic, roadworthiness or safety-performance claim. Specific legal and information-handling obligations need assessment for the business and its proposed use.

- [Car services and repairs](#), Consumer Affairs Victoria; current guidance checked 3 October 2026.
- [Guidance on privacy and the use of commercially available AI products](#), OAIC; published 21 October 2024, updated 17 January 2025, checked 3 October 2026. Refresh before live use.

If you are already trying AI and are unsure what deserves further investment, [tell us what you have tried and what you want to improve](#). We can help find the starting point together.

Read the structured HTML edition: alongsidepeople.com/resources/ai-for-automotive

INSURANCE

Improve claims service: give people a clearer file sooner

Alongside People / Advice and hands-on implementation

A practical AI workflow paper for Melbourne insurance operations.

Executive takeaway

If your team is already experimenting with AI, the next question may be which uses deserve confidence and further investment. We help make that decision clearer. You can start with what your people are trying and what concerns you; we can identify a suitable workflow together.

A claims colleague can receive a file containing the relevant information and still spend time reconstructing what happened: which document arrived when, what was already requested and which question remains unresolved. That preparation can crowd out the attention a customer needs.

AI may help organise an authorised file into a source-linked chronology and draft a factual status summary. The intended outcome is a clearer handover and less repeated administration. Coverage, settlement, underwriting and adverse decisions remain with authorised people. Preparation can influence judgement, so human checking must be substantive rather than a final click.

Our team helps design and implement bounded workflow changes, then test whether they improve the whole job. Start with staff support and evidence, not an unattended decision system. This paper offers a planning method; it reports no insurer results or guaranteed savings.

Find friction in the handover

Investigate a specific handover. Does a handler reopen correspondence to establish what was received? Does the next colleague ask a customer for information already supplied? Are supplier updates copied between systems? Does a manager have to rebuild the record before understanding why a file is waiting?

These are possible questions for discovery, not findings about all insurers or Melbourne teams. Follow the work with handlers, service staff and the people responsible for quality and conduct. Separate missing information from information that exists but is hard to find. Look for unclear ownership or inconsistent naming before deciding a model is needed.

Choose a narrow preparation task with an authorised reviewer. An easier handover should support customer service without classifying people as risky, inferring dishonesty or creating reasons to refuse a claim. A better document index or simpler status process may be sufficient.

A fictional running example

Fictional example, not a client case study: A Melbourne-based insurance operations team prepares an internal handover for a property claim. The file includes incoming documents, prior requests and supplier status updates. A handler checks a draft chronology against the original records and confirms the next administrative step.

The proposed tool lists received documents, source-linked dates and unresolved factual questions. It does not interpret coverage, infer fraud, recommend refusal, assess vulnerability, choose a settlement amount or contact the customer. It has no authority to change the claims record or trigger a payment. Questions that could affect a financial decision go to the authorised handler.

Testing begins with synthetic files. Real information requires the organisation's approved environment, access and use decisions. The example has no measured outcome and supplies no evidence of compliance.

What would make this worth continuing?

The following is a decision aid for the fictional trial, not an observed result. Agree the limits with the people responsible before testing.

- **Accepted result:** A handler-accepted factual chronology and handover with dates tied to original records.
- **Who judges it:** The receiving handler and the operational quality owner.
- **Hold and return to the manual route when:** A receipt is invented, a contradiction disappears or the draft makes an adverse or financial recommendation.
- **Evidence for another step:** Compare total preparer and receiving-handler effort for files of similar complexity; include repeated requests and missing facts.

A fictional chronology might show “Supplier email: proposed inspection Friday” beside “Handler note: inspection date not confirmed.” Keep both dated references and flag the conflict. The assistant must not silently turn the proposal into a confirmed booking.

Use the [one-workflow worksheet](#) to record the agreed question, evidence and next decision.

Five stages with proof before expansion

Understand. Enter with one handover task, its outcome owner and the people receiving the file. Map receipt, indexing, preparation, checking and next-step ownership. The benefit hypothesis is less reconstruction and fewer repeated information requests. Exit with a baseline, authoritative source hierarchy, permission boundaries and a clear definition of an acceptable handover. If staff disagree about what a source means, resolve it through the existing process rather than asking AI to decide.

Try. Enter after scope, reviewers, permitted data and stop conditions are agreed. Use synthetic files with duplicate documents, conflicting dates and a supplier update that contradicts an earlier note. The hypothesis is quicker factual preparation without losing uncertainty. Exit with reviewed observations,

correction effort and the owner's decision. Stop if the draft invents a receipt, suppresses a contradiction or recommends an adverse decision. Restore the manual handover.

Repeat. Enter when the trial provides a reason to continue. Standardise references, event labels and how gaps are recorded. Have a second colleague follow the process and compare each draft with its sources. The hypothesis is more consistent handovers between people. Exit with reproducible instructions, exception records, a manual fallback and comparison against the baseline. Include awkward files; excluding them can conceal the burden the process creates.

Operate. Enter only after bounded routine use is approved. Name the owner and backup, control record access and make failed preparation visible. A handler must know when a chronology is incomplete and how to reconstruct it. The hypothesis is steadier service and fewer internal chases. Exit with operating guidance, change controls, visible exceptions, exercised recovery and an owner review. Pause when permissions, source reliability or the ability to recover changes.

Improve. Enter with enough observations to evaluate usefulness. Compare total handling effort, handover quality and customer-facing consequences. The hypothesis is keeping helpful preparation while preventing inappropriate expansion. Exit each review with a recorded retain, revise, reduce or retire decision and next trigger. If a structured event log proves better than generated prose, reduce AI use. New financial or customer actions require a separately authorised design and evidence decision.

This **Understand / Try / Repeat / Operate / Improve** framework is Alongside People's original per-workflow discussion aid. It is not an official insurance standard, certification or validated readiness test. Greater autonomy is not a requirement for progress.

Respect the regulatory and authority boundaries

ASIC's [claims-handling information sheet, INFO 253](#) explains the AFS licensing regime and obligations for claims handling and settling services involving insurance products regulated by ASIC. The relevant licensee and authorised specialists must determine how a proposed workflow fits those obligations.

For APRA-regulated entities within its scope, the current [CPS 230 Operational Risk Management](#), whose current revision commenced on 1 July 2026, addresses operational risk, critical-operation continuity and service-provider risk. Do not assume every intermediary is an APRA-regulated entity or that a supplier arrangement has a particular classification without review.

OAIC's [commercial AI privacy guidance](#) addresses personal-information risks and product due diligence. Insurance files may require careful handling; determine actual obligations and approved controls with the organisation's specialists. This paper gives no legal assurance, licence coverage or compliance sign-off.

A proposed 30-day pilot

This schedule is a planning example, not our service promise. Begin after the sponsor has approved the people, scope and environment. A calendar target does not bypass information, conduct or procurement review.

Days 1-7: Follow existing handovers and record preparation, review and rework. Agree a synthetic example set, source hierarchy and acceptance criteria. Explicitly exclude coverage, fraud, underwriting, settlement and other adverse recommendations.

Days 8-14: Build a draft chronology and factual status summary. Handlers check each event against its source. Test contradictions, absent documents and poor document quality. Record whether checking effort offsets any preparation benefit.

Days 15-23: Repeat in shadow mode with authorised examples. Existing records remain authoritative. Test a second handler's use, an unavailable tool and recovery through the manual route. Nothing is sent externally or used to automate a financial decision.

Days 24-30: Review outcomes and exceptions with the operational owner and appropriate risk/conduct stakeholders. Deliver the evidence, limitations and operating proposal. Choose further testing, a separately approved next step, revision or stopping. Routine-use readiness is an explicit decision.

Measure service, not just generated output

The primary outcome could be staff effort to prepare an accepted handover. Count indexing, drafting, source checking, correction and the receiving handler's reconstruction. Record elapsed handover turnaround separately. Pair these with repeated information requests, unresolved discrepancies and usefulness to the next colleague.

Guardrails include incorrect dates, missing contradictions, unsupported assertions, unauthorised access and adverse recommendations. Review whether the process adds work for customers or suppliers. Any escalation related to customer needs stays with the existing authorised process; the pilot does not classify vulnerability or treat a model's inference as fact.

Keep criteria fixed, compare similar file complexity and document sample limits. A shorter internal task does not automatically establish faster settlement or better customer outcomes. Verify amounts through established records and controls; fluent text is not financial validation.

Before implementation, record supplier data access, retention, processing locations, subcontractors and exit arrangements. Limit access to necessary information and preserve authoritative records. Test what happens when the tool is unavailable, so a preparation aid does not become a hidden dependency that obstructs service.

How our team helps

Alongside People provides practical AI advice and hands-on implementation. Paul Volpato, our founder and CEO, connects the purpose of the work with the people, systems and checks it needs. Our team can map the handover, design the pilot, implement an agreed preparation workflow and develop operating instructions, review checks and handover material.

Your authorised professionals retain financial and customer decisions. We agree scope, responsibilities and support arrangements rather than imply insurance credentials or indefinite operation. Bring a recurring handover, the person receiving it and a clear service problem. We can help make the next change concrete enough to assess.

Official references

Accessed 3 October 2026. Recheck before consequential application.

- [ASIC INFO 253](#): regulator information sheet, scoped to ASIC-regulated insurance claims handling.
- [APRA CPS 230](#), current revision commenced 1 July 2026: prudential standard for its defined APRA-regulated entities; supersedes the previous version.
- [OAIC commercial AI privacy guidance](#): privacy guidance; does not cover every applicable regime.

If you are already trying AI and are unsure what deserves further investment, [tell us what you have tried and what you want to improve](#). We can help find the starting point together.

Read the structured HTML edition: alongsidepeople.com/resources/ai-for-insurance

CREATIVE

Make room for the good idea

Alongside People / Advice and hands-on implementation

Practical AI adoption for Melbourne creative teams

If your studio is already trying AI but has no clear view of what good use looks like, we can help you work it out. We look at what supports the quality of the work, what adds burden and where to go next. You do not need to arrive with a production task already chosen.

A useful creative workflow leaves people enough attention to make and defend good decisions. It makes the brief clear, keeps feedback attached to the right version and carries an approved idea through production without losing its intent.

Our starting point is to look first at the repeated work surrounding the idea. AI can be trialled for organising a brief, preparing production variations or checking a delivery packet. The creative lead still decides what deserves to exist, what fits the audience and whether the final work is good enough to represent the client.

For Melbourne creative teams, we propose **Understand, Try, Repeat, Operate, Improve**. Protect judgement while reducing production friction. Generated options need to earn their place.

Locate the work that steals attention

A project can slow down before anyone opens a design tool. The brief contains several competing objectives. Feedback arrives in separate messages. A late amendment changes the product claim but reaches only some files. The approved concept is clear, yet turning it into a set of formats requires repeated copying and checking.

Brief clarification needs a conversation. Version confusion needs a reliable record. Repeated production may benefit from templates or automation. Assistance depends on clear purpose and sources.

Producing many directions can make review harder. Decide what an exploration needs to answer and how many alternatives the team can assess. Include creative-lead attention in the cost.

Begin by following one recent deliverable from briefing to acceptance. Record where people searched, repeated, waited or corrected. Ask which decisions required creative judgement and which steps simply carried an agreed decision forward.

A fictional workflow: one campaign, several formats

This fictional example reports no client results or achieved savings.

Imagine a Melbourne studio preparing a campaign for a local cultural event. The client has approved the audience, message, visual direction and a set of supplied photographs. The studio needs to prepare several digital formats and a presentation explaining the choices.

The producer creates an approved source packet: brief, required wording, factual event details, brand rules, permitted assets, output dimensions and the client's requirements for AI use. An assistant organises these into a production checklist. It flags missing information rather than inventing an admission price, partner endorsement or usage permission.

The creative lead reviews the checklist and chooses which tasks can be assisted. Layout exploration uses placeholders or assets cleared for that purpose. The designer determines composition, hierarchy and tone. Where generative imagery is proposed, its provenance and permitted use are checked before it becomes part of a client deliverable.

The assistant helps prepare draft text variations within the approved message. Each variation remains attached to its source wording and format. The producer checks dates, names, restrictions and file naming. The creative lead reviews whether each version preserves the idea, rather than merely filling the available space.

Client feedback is recorded against a specific version. An approved amendment becomes a change instruction for the affected outputs. A final packet contains editable files, exports, approval history and an asset register. A person signs off the actual delivered files.

The trial examines chasing, correction and creative quality. A folder of variants alone is insufficient evidence.

Rights belong in the workflow

Tool capability, permission to use an input, contractual commitments and rights in an output are separate questions. A client's supplied photograph may be cleared for a campaign without being cleared for every upload or transformation. Record the allowed use before sending material to a new service.

Adobe's Firefly FAQ describes commercial use of generated outputs subject to the actual feature and its terms. Check the feature, model, input permissions and account arrangement before production. A familiar application name does not settle the rights for every generation, especially where partner models are available. Keep the applicable terms with the asset record. [Firefly FAQ](#).

WIPO's current small-business guide explains that AI and intellectual-property questions are evolving and jurisdiction-dependent. We recommend a practical source and approval record, with specialist advice where the proposed use requires it. Do not promise exclusivity or legal clearance from the fact that a tool produced an image. [WIPO: Learning Machines](#).

What would make this worth continuing?

The following is a decision aid for the fictional trial, not an observed result. Agree the limits with the people responsible before testing.

- **Accepted result:** Accepted campaign files with the approved message, editable source and complete asset and approval records.
- **Who judges it:** The creative lead and producer, with the client's normal approval process.
- **Hold and return to the manual route when:** Asset permission is missing, a factual claim changes or an approved amendment is missed in an affected output.

- **Evidence for another step:** Compare accepted deliverables of similar scope, including unusable variants, creative review and correction rounds.

For a fictional amendment test, change the event date in the approved source packet. List every affected format and verify the actual exported files. An updated prompt or master file is insufficient if one delivered version still carries the old date.

Use the [one-workflow worksheet](#) to record the agreed question, evidence and next decision.

Understand: separate judgement from production

We map one deliverable with the producer, creative lead and people making the work. We identify the decisions that define quality and the repeated steps that can be made easier. We agree the source of truth for the brief, versions and approvals.

The aim is focused creative attention. Continue when there is an agreed workflow, a quality rubric, an asset-use record and a baseline. The rubric might cover message accuracy, brand fit, visual hierarchy and fidelity to the concept. It should be specific enough for reviewers to explain a rejection.

Try: carry one approved idea through a contained task

We select a task such as organising a source packet or preparing draft text variations for approved formats. The trial uses synthetic or authorised material. A designer and producer inspect the outputs and record the changes needed before acceptance.

The aim is evidence about useful assistance. Continue when there is a finished sample, editable source, review notes and full effort record. Record unusable outputs as well as accepted ones.

Repeat: test the method against real variation

We repeat with a second brief or a controlled amendment. Can the method preserve a changed date across every affected version? Does it retain the approved voice? Does another producer understand the record? Include a case with insufficient rights information so the workflow must stop and ask.

The aim is repeatable production. Continue when there is a reusable template, exception handling and comparable samples that preserve intent while allowing creative judgement.

Operate: make the approval trail easy to use

We establish where briefs, assets, versions and approvals live, who can access them and who can release work. Tool choices include their actual edition and model arrangements. The team can identify which files used assistance and recover editable work without relying on a chat history.

The aim is less uncertainty at handover. Continue when there is a working procedure, an asset register, a tested delivery packet and a manual fallback. Client confidentiality, permissions and publishing authority remain explicit. Automated production should not quietly become automated approval.

Improve: retain what helps the work

We review accepted quality, review burden and correction patterns. If output is faster but creative review becomes harder, narrow the task or reduce the number of options. If templates account for most of the improvement, keep that simpler method. If a tool changes, retest the parts that depend on it.

The aim is a production practice that earns its place. Continue when there is a dated decision and a focused next experiment.

A proposed 30-day trial

This proposed evaluation is not a commercial commitment or campaign-delivery promise.

During days 1-5, choose one deliverable, trace comparable recent work and agree quality and rights requirements. During days 6-12, complete a contained sample from a checked source packet. During days 13-21, repeat with another example and an amendment, inspecting every output. During days 22-30, review the evidence and decide whether to incorporate the method into normal production.

Keep existing client approvals. Fictional briefs can demonstrate the method but cannot establish client acceptance. Record the limits and next approved observation.

Measure accepted work, including the checking

Establish a baseline for comparable deliverables. Record elapsed brief-to-acceptance time, active production time, producer and creative-lead review time, correction rounds and missing or inconsistent requirements. Separate a change in client direction from a production mistake.

Assess quality with the agreed rubric and written reasons. A review of matched examples can be useful, provided reviewers know the intended audience and message. Small samples should inform a decision without being presented as statistical proof.

Count accepted deliverables with a complete approval and asset record. Track factual corrections, unapproved assets and missed amendments as quality guardrails. Pause the workflow if required provenance or permission cannot be established. Do not reward volume of imagery, prompts or variants.

Calculate cost per accepted deliverable from tool usage, production, review and rework. Show initial setup separately. Record whether producers can find the latest approved version and whether creative leads have more room for the decisions that matter. These observations explain a result better than a headline drafting-time saving.

How we help

Alongside People offers advice and implementation around the work. We can help clarify the production problem, design a usable source and approval packet, assess tools against the task, and build a repeatable method. We help your team measure the result and decide how far to take it.

Paul Volpato is our founder and CEO. Start with one deliverable where the team spent too much attention finding, resizing, updating or checking.

Sources and scope

Understand / Try / Repeat / Operate / Improve is Alongside People's original discussion framework. It is not a certification, validated maturity assessment or official standard.

The five stages and trial are our proposed approach. The workflow is fictional and makes no claim about client outcomes. Sources support specific product and intellectual-property context; they do not establish rights for a particular asset or project.

- [Adobe Firefly FAQ](#), Adobe; displayed update 25 May 2026, checked 3 October 2026.
- [Learning Machines: introduction](#), World Intellectual Property Organization; 2026 small-business guide, checked 3 October 2026.

If you are already trying AI and are unsure what deserves further investment, [tell us what you have tried and what you want to improve](#). We can help find the starting point together.

Read the structured HTML edition: alongsidepeople.com/resources/ai-for-creative

HEALTHCARE

More time for care: improve the administration around it

Alongside People / Advice and hands-on implementation

A practical AI paper for Melbourne healthcare organisations.

If staff are already trying AI but its useful place in your organisation remains unclear, we can help you make sense of the options. Together we establish what better would look like and where to begin. The administrative example below is one way to explore that question.

The useful starting point. Start with the coordination around care: finding the right procedure, preparing an internal brief and keeping administrative instructions current. A useful improvement gives staff capacity for more valuable work while preserving privacy, accuracy and accountability. The first test is whether the whole job becomes easier, including checking and correction.

Where the opportunity begins

A healthcare administrator may spend a morning reconciling contradictory instructions before answering a straightforward operational question. Which contact owns a facilities request? Which booking procedure is current? What needs to be included in an internal service update? Finding and checking the answer can take more effort than writing it.

These are useful situations to examine with staff. We should establish whether the friction comes from inaccessible information, unclear ownership, outdated documents or repetitive preparation. AI might help assemble a source-linked draft; a better document register or simpler form may resolve more of the problem.

Our recommendation is to start away from clinical decisions. This paper concerns internal administrative preparation. Diagnosis, triage, treatment recommendations, patient prioritisation and decisions that restrict access to care are outside its proposed pilot. We also exclude clinical notes and ambient recording. Those uses require their own specialist assessment and authority.

Administrative work still deserves careful boundaries. A roster, correspondence thread or appointment record can reveal personal or health information. Calling a task “back office” does not make its inputs suitable for an AI tool.

The Victorian context

The Victorian Department of Health now provides minimum standards, governance models and escalation material for AI in Victorian public health services. Its source page describes requirements that must be met before use, rather than treating an experiment as an exemption. The documents were updated on 3 June 2026; the page was updated on 7 July 2026. These public-health requirements should not be presented as a universal policy for every private practice. [Victorian public health AI resources](#).

The Department's explanation of the Health Records Act says its Health Privacy Principles apply to health information handled in Victoria across public and private sectors. Identifying information collected while providing health services may itself be health information. The practical implication is to ask the organisation's information and privacy owners about the actual data, not assume a spreadsheet is harmless because it lacks a diagnosis column. [Health Records Act explanation](#).

For organisations within its scope, OVIC's enterprise AI guidance expects preparation that includes privacy impact assessment, information-security assessment, training, incident response and continuing review. Buying an enterprise tool does not settle those questions. [OVIC enterprise AI guidance](#).

These sources support planning; this paper does not determine legal compliance or approve a tool.

A fictional workflow: prepare the weekly administration brief

A fictional Melbourne health service circulates a weekly brief about facilities contacts, approved administrative procedures and operational notices. A coordinator gathers the material; department managers check it before circulation. Several drafts repeat old instructions because a superseded notice remains easy to find.

The proposed change uses a small approved collection of non-personal administrative documents. AI prepares a draft identifying each statement's source, document version and unresolved conflict. It can report that two instructions differ. It cannot decide which is authoritative, invent a contact, access patient systems or send the brief.

The document owner resolves conflicts. The coordinator checks completeness and the manager approves circulation. A visible "not confirmed" item is preferable to a fluent answer that hides missing evidence. This is a constructed example, not an account of our healthcare work.

What would make this worth continuing?

The following is a decision aid for the fictional trial, not an observed result. Agree the limits with the people responsible before testing.

- **Accepted result:** A manager-approved administration brief with current sources and unresolved conflicts made visible.
- **Who judges it:** The coordinator, document owners and managers who use the brief.
- **Hold and return to the manual route when:** A patient source is accessed, personal information appears or an incorrect instruction could affect a consequential action.

- **Evidence for another step:** Compare accepted briefs of similar scope, including source verification, correction and manager effort.

A fictional permitted-source example is a non-personal facilities notice identifying an approved request procedure. An appointment list is different: it may reveal health information even without a diagnosis column. The organisation must decide which real sources are permitted.

Use the [one-workflow worksheet](#) to record the agreed question, evidence and next decision.

Five states, with evidence before expansion

Understand. Follow one brief from request to approved circulation. Ask the recipient which missing details cause follow-up. Identify source owners, permitted documents and approvals; record preparation and review effort separately. The potential benefit is less searching and fewer conflicting instructions. The gate is a named owner, a clear problem, authorised information and a baseline. If documents cannot be classified or approved, stop before connecting a tool.

Try. Compare a manual brief with a draft prepared from synthetic or approved non-personal examples. Include an expired notice, a missing date and two contradictory instructions. The potential benefit is faster assembly. The gate is traceable facts, visible uncertainty and checking effort within the agreed limit. No generated claim passes merely because its wording sounds professional.

Repeat. Have a second coordinator use the documented instructions on a different set of notices. Record corrections and reasons for manual handling. The potential benefit is dependable preparation across staff changes. The gate is comparable quality across representative repetitions, including an awkward example, and a fallback that someone other than the author can follow.

Operate. Only with explicit approval, establish routine drafting within the agreed collection. Name an owner and backup, limit access, retain source versions and test a pause. The potential benefit is less coordination effort around a recurring brief. The gate is exercised recovery, visible exceptions and acceptance of the ongoing workload. Routine use does not authorise access to patient records.

Improve. Review whether staff actually have more useful capacity and whether recipients trust and use the brief. Retire duplicated notices or improve source ownership if those changes matter more. The gate is a recorded retain, revise or stop decision supported by observations. More AI autonomy is not a required next step.

This original Alongside People framework is a discussion aid, not a health-service accreditation, validated maturity score or official standard.

A proposed 30-day pilot

Days 1-7: follow the current process, select the document collection, obtain relevant approvals and record a baseline. Days 8-14: prepare paired drafts offline with agreed examples; inspect every output. Days 15-23: repeat with another staff member and exercise missing-source and tool-failure scenarios. Days 24-30: review evidence, operating burden and the next decision with the sponsor.

If approvals or source preparation take longer, the timetable changes. Live circulation is not necessary to learn from the pilot. The proposal is a planning example, not a promise of a 30-day engagement or permission to use real health information.

Measure useful capacity

Choose accepted-brief preparation effort as the primary measure, including searching, prompting, review and correction. Record elapsed time separately. Preserve the original baseline and compare briefs of similar scope; show differences rather than attributing all change to AI.

Guardrails include unsupported statements, wrong or expired instructions, missing sources, review effort, inaccessible wording and any inappropriate information exposure. Include tool charges, configuration, training and ongoing document maintenance in cost. Agree tolerances before trialling; an information-handling incident or consequential incorrect instruction triggers a pause and the organisation's incident process.

Recovered minutes are potential capacity, not automatically cash savings or additional patient appointments. Ask staff what useful work became possible and whether the change shifted effort onto managers.

Permissions, recovery and our role

Grant access only to the approved collection. Retain original documents and drafts in authorised systems, use the organisation's retention rules and treat retrieved content as information rather than instructions to the automation. No sending, record amendment or source deletion belongs in the initial tool permissions.

On failure, hold the draft, identify affected material and return to the established manual brief. Reconcile any previously circulated error through the authorised correction process. A backup coordinator should complete this exercise before routine use.

Our team can help understand the work, design the experiment and implement the agreed drafting and review process. We can connect approved information, document responsibilities and prepare a usable handover. Your organisation retains clinical, information-governance and operational authority; specialist approval and ongoing support are separately agreed.

Bring us a recurring administrative job, the people doing it and a representative example that can be shared. We can help establish whether a practical change is worth making.

Source check: 3 October 2026. Official sources above support the narrow policy statements; the workflow, five states and pilot are our proposed method. Recheck applicable requirements before use.

If you are already trying AI and are unsure what deserves further investment, [tell us what you have tried and what you want to improve](#). We can help find the starting point together.

Read the structured HTML edition: alongsidepeople.com/resources/ai-for-healthcare

EDUCATION

Give education staff more time for the work that needs them

Alongside People / Advice and hands-on implementation

A practical AI paper for Melbourne education organisations.

You may already have staff exploring AI while still working out what good use should look like across your institution. We help you turn those experiences into a clearer direction and choose a worthwhile next step. There is no need to have an improvement project defined first.

The useful starting point. Improve recurring staff administration. An accurate internal operations pack from approved material is a useful first experiment. Test whether it reduces searching, rewriting and correction, leaving capacity for work requiring staff attention and judgement.

Start with the staffroom's recurring work

A useful question for an education leader is what staff repeatedly have to reconstruct. The answer might be room-use instructions, event preparation steps, a meeting pack or an induction guide. Instructions may be spread between a policy library, an email thread and a document someone keeps updating privately.

An administrator may not know which checklist is current or who must approve a change. Asking AI to rewrite everything risks making those gaps less visible.

Follow one task with its producers and users. Establish an acceptable result, authoritative information and human-owned decisions before choosing AI or ordinary document management.

This paper focuses on staff administration. Student ranking, grading, admission, disciplinary decisions, welfare assessments and individualised learning judgements are outside the proposed pilot. Student reports and direct communication with families are also excluded. The people and relationships at the centre of education remain a reason to design the process carefully.

Know which guidance applies

The Victorian Department of Education's generative AI policy sets requirements for Victorian government schools, including restrictions on uploading personal and sensitive information. It directs staff away from AI use that replaces teachers' and leaders' professional judgement or authentic communication. Its requirements must be read with other departmental technology and privacy policies. These school rules should not be described as the policy of every university, TAFE or independent provider. [Victorian school AI policy](#).

The accompanying appropriate-use guidance recommends disclosure, records and monitoring, and asks schools to consider administrative and financial costs. That supports checking the whole workflow rather than celebrating faster text generation. [Victorian appropriate-use guidance](#).

The Australian Framework for Generative AI in Schools guides responsible and ethical use for people connected with school education. Education Ministers endorsed its 2024 review in June 2025. It is a national guidance framework, not approval for a particular product or a replacement for institutional requirements. [Australian schools framework](#).

Before implementing a generative AI tool, the Victorian policy also requires schools to ensure accessible and inclusive use. Include the relevant reasonable adjustments in the proposed workflow and trial, rather than leaving accessibility to the final export.

A tertiary provider should identify its own approved-tool, information, accessibility and records requirements before adopting this example. We offer a practical planning method, not a determination of compliance.

A fictional workflow: the staff induction pack

A fictional Melbourne secondary school prepares an induction pack for new staff each term. The administrator gathers public or internally approved non-personal instructions about facilities, document locations and routine processes. A leader checks the pack. Old copies contain inconsistent room names and links to superseded procedures.

The proposed assistant drafts sections from a deliberately small source collection and attaches a link and version to every instruction. It produces a separate list of missing or conflicting information. It cannot read student records, analyse staff performance, answer welfare questions or distribute the pack.

The administrator checks the links and facts; process owners confirm the instructions; a leader approves the result. A new staff member then uses the draft to find an ordinary procedure and reports what remained unclear. This tests usefulness as well as correctness.

Use role names rather than a personal staff directory. This fictional example does not imply previous school engagements.

What would make this worth continuing?

The following is a decision aid for the fictional trial, not an observed result. Agree the limits with the people responsible before testing.

- **Accepted result:** An approved induction pack that a new staff member can use to find the intended procedure.
- **Who judges it:** The content owner and approving leader, with a new staff member testing usefulness.
- **Hold and return to the manual route when:** Personal information enters the trial, an instruction lacks a source or an omitted step could misdirect a consequential action.
- **Evidence for another step:** Compare total preparation and clarification effort for similar packs; test source changes and accessible use.

In a fictional check, ask a new colleague to find the current room-booking procedure using the pack. Record whether the link works, the instruction is clear and the format is usable with their access needs. The trial has more to learn if the colleague still needs to find the author.

Use the [one-workflow worksheet](#) to record the agreed question, evidence and next decision.

Five states that keep the decision practical

Understand. Follow the pack's preparation and a new staff member's attempt to use it. Identify duplicated instructions, authoritative owners and approval steps. The benefit hypothesis is less reconstruction and fewer follow-up questions. The evidence gate is an agreed scope, approved sources, named owner and baseline. If the real issue is a missing policy decision, resolve that first.

Try. Prepare an offline draft from approved non-personal or synthetic material. Include a broken link, contradictory room instructions and a missing approval step. The benefit hypothesis is less assembly time. The gate is a draft with traceable instructions, visible gaps and manageable review effort. Student or staff personal information entering the trial triggers a pause.

Repeat. Have another administrator prepare a new version using the same documented process. Include a changed source. The benefit hypothesis is consistency between terms and colleagues. The gate is reproducible quality, correct handling of changes and a usable manual route. A checklist that works only for its author needs further refinement.

Operate. Introduce the approved process within a defined scope, with a content owner, backup and review calendar. Keep dissemination under human control. The benefit hypothesis is less chasing for a dependable pack. The gate is an exercised fallback, appropriate access and an accepted maintenance workload. An extra feature in the software does not expand permission.

Improve. Review actual preparation effort and the recipient's ability to find and act on instructions. Remove redundant content and correct the source library before adding more generation. The benefit hypothesis is continuing usefulness rather than increasing volume. The gate is a documented decision to retain, revise or stop, with evidence and a next review trigger.

Understand / Try / Repeat / Operate / Improve is our original discussion framework. It is not a validated assessment, school quality rating or official standard. Different workflows may be in different states.

A proposed 30-day pilot

In days 1-7, select one pack, observe the manual workflow and confirm source and tool approval. In days 8-14, compare manual and assisted drafts using the same material. In days 15-23, repeat with another preparer, check an updated procedure and run a short usability exercise. In days 24-30, review the results and decide whether routine use is justified.

Schedule the work around staff capacity and the institution's approval process. A month without an appropriate event can still support a simulation, but it cannot establish performance during a busy term. These are proposed steps, not a commercial timetable, promised result or live-use authorisation.

Count work returned to staff

The primary measure is total staff effort to produce an accepted pack. Include preparation, review, fixing links and answering clarification questions. Compare documents of similar scope and record whether policy changes altered the workload.

Quality guardrails include incorrect instructions, omitted approval steps, outdated sources and whether a recipient can complete the intended task. Check accessible formatting and clear language with actual users. Record review time separately: faster drafting that requires longer checking may not help.

Cost includes licences, initial setup, training and maintaining the collection. Agree acceptance thresholds before the test. Stop if personal information appears, if an instruction could misdirect a consequential action or if the process cannot identify its source. Preserve unsuccessful drafts so errors inform the decision.

Additional capacity is a benefit hypothesis. Ask staff what they could spend time on instead; do not infer improved student outcomes or financial savings from a quicker induction document.

A process people can maintain

Limit permissions to the approved source collection and a draft workspace. Keep versions, responsible owners and approval status visible. Documents retrieved by the assistant must not instruct it to send messages or broaden access. Distribution and source amendments remain human actions in the pilot.

If drafting fails, return to the existing template and document list. Hold unverified sections rather than silently substituting a plausible answer. If an approved pack contains an error, use the school's correction route and identify which copies need replacing. Test that a backup administrator can prepare the pack manually.

Our team can map the work, design a bounded trial and implement the agreed process with staff. We can help organise approved information, connect suitable tools and produce instructions, checks and a handover. Institutional leaders retain policy, educational and publication decisions. Wider rollout and continuing support require separate agreement.

Sources inspected 3 October 2026. The linked policy and guidance support the specific statements above; refresh applicable institutional requirements before implementation.

If you are already trying AI and are unsure what deserves further investment, [tell us what you have tried and what you want to improve](#). We can help find the starting point together.

Read the structured HTML edition: alongsidepeople.com/resources/ai-for-education

MANUFACTURING

Less chasing, clearer production handovers

Alongside People / Advice and hands-on implementation

A practical AI paper for Melbourne manufacturers.

If AI experiments are happening around the business but you are unsure which ones deserve investment, we can help establish what good looks like for your operation. We work with your people to find a useful next move; the handover below illustrates one possibility.

The useful starting point. Recover capacity by making production documentation easier to assemble and verify. Start with a handover or review pack that carries the right facts, sources and unresolved questions to the next person. Keep safety, quality release and production authority with the people who hold it.

The work between the systems

A production planner may have the order in one system, an approved drawing in another and the latest supplier clarification in an email. Preparing a handover requires finding the documents, checking their revisions and asking whether an unresolved issue has actually been resolved. A neatly written summary can still be wrong if it joins information from different versions.

Investigate these possible bottlenecks through one job's planning and handover. Ask the receiving supervisor which facts need chasing. Observe how engineering changes, partial information and urgency affect preparation.

AI might prepare a source-linked pack and identify gaps. A document register or required fields may help more. Choose the approach that makes information dependable with less preparation.

This paper proposes administrative documentation support. It does not propose autonomous machine control, altered safety settings, generated operating instructions, maintenance authorisation or product release. An assistant must not convert a suggestion into an approved technical decision.

Safety duties remain part of the setting

WorkSafe Victoria describes employer duties to provide safe plant, safe systems of work, and necessary information, instruction, training and supervision. Delegating document preparation does not transfer those duties to a tool. [WorkSafe general duties](#).

WorkSafe's consultation guidance explains consultation with affected employees and health and safety representatives. Changes to a handover may change how people receive information or organise work, so involve those users and the organisation's safety owners rather than designing solely for the planner. The applicable duty and process should be checked for the actual change. [Consultation: safety basics](#).

The Plant compliance code offers practical guidance for relevant Victorian duties and is not itself mandatory. Check the controlled source and its applicability to the actual work. This paper does not use the code to design or certify machinery. [Plant compliance code](#).

These are official Victorian resources. Their presence in a document does not establish legal compliance, engineering verification or safety approval.

A fictional workflow: prepare the job review pack

A fictional Melbourne component manufacturer holds a planning review before starting an order. Its planner gathers the approved drawing reference, purchase-order details, material-document references and open supplier questions. Engineering and quality staff verify the pack before the supervisor accepts the handover.

The proposed assistant reads a restricted set of approved records and prepares a draft index and factual summary. Each item names its source and revision. It flags mismatched order references, missing attachments and contradictory dates. It does not infer that a material is suitable or that a deviation has been accepted.

A document controller checks source status. Engineering resolves technical questions; quality retains acceptance authority; the planner completes the administrative handover. The initial trial uses synthetic orders and documents, including deliberately stale versions.

The pilot never changes the production schedule, bill of materials, equipment parameters or inspection result. It cannot create a compliance certificate. A missing approval remains missing. This fictional example describes a design to test, not our delivery history.

What would make this worth continuing?

The following is a decision aid for the fictional trial, not an observed result. Agree the limits with the people responsible before testing.

- **Accepted result:** An accepted job review pack with correct identifiers, source revisions and missing approvals visible.
- **Who judges it:** Planning, document control and the engineering or quality staff receiving the pack.
- **Hold and return to the manual route when:** An approval is invented, a superseded drawing is presented as current or the draft issues a technical or safety instruction.
- **Evidence for another step:** Compare total preparation and verification effort for similar jobs, including corrections and supervisor questions.

In a fictional revision check, the index contains drawing revision B while the controlled register identifies revision C as current. Hold the affected item and ask document control to resolve it. The assistant cannot merge the drawings or infer that the difference is harmless.

Use the [one-workflow worksheet](#) to record the agreed question, evidence and next decision.

Five states with manufacturing evidence gates

Understand. Observe one pack being assembled and used. Separate search and formatting effort from engineering and quality judgement. Record authoritative sources, document ownership and the manual checks. The benefit hypothesis is less searching and fewer incomplete handovers. The gate is a named owner, agreed scope, baseline and permission to use the chosen records. If revision status is unclear, fix that before adding AI.

Try. Draft a pack from synthetic or explicitly approved examples. Include a superseded drawing, mismatched part number and absent approval. The benefit hypothesis is quicker factual assembly. The gate is complete source traceability, correct identifiers and revision handling, visible gaps and acceptable review effort. One plausible invented approval is grounds to stop and correct the design.

Repeat. Have a different planner use the instructions across more representative order types. Include incomplete supplier material. The benefit hypothesis is consistent handovers between people and shifts. The gate is repeatable performance without silently dropping hard cases, plus a manual process that the receiving team understands. Record exclusions openly rather than reporting only easy jobs.

Operate. After authorisation, connect the limited documentation workflow with access controls, logging, a backup owner and change review. The benefit hypothesis is less coordination effort in routine preparation. The gate is demonstrated recovery and explicit acceptance of the ongoing control workload. A drafting process remains outside machine control and quality-release authority.

Improve. Review errors, delays, user experience and actual capacity recovered. A better required field or clearer document ownership may outperform a more capable model. The gate is a recorded decision to retain, revise, expand or retire, with engineering and safety review where the next change requires it. Expansion needs fresh permission and evidence.

Our Understand / Try / Repeat / Operate / Improve framework is an original discussion aid. It is not a manufacturing certification, validated scoring system or substitute for established quality and safety practices.

A proposed 30-day pilot

Days 1-7: map one documentation handover, identify document owners and record a baseline. Confirm information permission and technical boundaries. Days 8-14: prepare paired manual and assisted packs offline, then review every identifier and source. Days 15-23: repeat with another planner and exercise stale-source, missing-attachment and tool-outage cases. Days 24-30: review results with planning, engineering, quality and affected users.

The decision may be to improve document control before continuing. The month is a suggested learning period, not a delivery promise or justification to skip approval. It cannot establish performance for all products, shifts or abnormal production conditions.

Measure the accepted handover

Use total preparation and verification effort per accepted pack as the primary measure. Separate active effort from waiting for a supplier or engineer. Record first-pass completeness, follow-up questions and reasons for rejection. Compare similar jobs and retain the original manual examples.

Guardrails include wrong identifiers, wrong revisions, omitted approvals, unsupported technical statements and inspection burden. Confirm that the pack's layout makes uncertainty easy to see. A reduction in planner effort must not create greater checking effort for engineering or the supervisor.

Include licences, configuration, integration, training and document maintenance in cost. Agree thresholds before the pilot. Any invented safety instruction, approval or release statement triggers a hold and review; the affected job continues under established manual controls. Do not infer reduced scrap, improved output or shorter lead times without evidence from those outcomes.

Recovered capacity may help an unresolved order. It does not justify faster production or assumed labour savings.

Permissions, fallback and our role

Begin with read access to the selected document set and write access only to a draft workspace. Keep original records immutable to the assistant. It cannot approve deviations, place orders, edit controlled documents or send instructions to machines or workers.

If a source becomes unavailable or revisions conflict, hold the pack and identify the item needing human resolution. Return to the existing manual checklist. Before routine use, ask a backup planner to reconstruct a pack from the originals and reconcile a failed run without duplicating records.

Our team can help follow the work, assess alternatives and implement the agreed documentation improvement. We can connect authorised information, make evidence and uncertainty visible, define checks and prepare a handover. Your organisation retains engineering, quality, safety and operational authority. Production-system changes and ongoing support need separately agreed scope.

Bring us a handover that repeatedly causes chasing or rework. We can help establish whether better information and a bounded assistant would make the work easier to rely on.

Sources inspected 3 October 2026. Recheck source versions and applicable requirements before use.

If you are already trying AI and are unsure what deserves further investment, [tell us what you have tried and what you want to improve](#). We can help find the starting point together.

Read the structured HTML edition: alongsidepeople.com/resources/ai-for-manufacturing

LOGISTICS

Resolve delivery exceptions with less chasing

Alongside People / Advice and hands-on implementation

A practical AI paper for Melbourne logistics teams.

Your team may already be using AI in small ways and still be unsure whether it is making the operation better. We help you assess what is happening, define a useful result and identify the next opportunity. The exception below is an illustration, not a project you need to have selected in advance.

The useful starting point. Give coordinators a clearer view of an exception: what is confirmed, what is missing and who can decide the next step. AI may help assemble that picture from authorised records. The useful outcome is less administrative chasing and more capacity for responsible coordination, with dispatch and safety decisions remaining human-owned.

Follow the exception, not just the delivery

An exception can create a second job around the transport task. A delivery reference differs between an email and a warehouse record. A booking changes while one team is still working from the earlier notice. Customer service asks for an update before the coordinator has confirmed what happened.

Investigate where finding facts and contacting the right person causes delay. An AI summary helps only if it distinguishes confirmed events from estimates and ties statements to sources.

A draft collects references and unresolved questions. A coordinator verifies the situation and obtains required dispatch, customer or facility decisions.

The pilot excludes autonomous driver or vehicle allocation, route changes, fatigue assessments, load decisions and work/rest instructions. It also excludes worker ranking and messages that pressure a driver to meet an unsafe target. A shorter message queue is useful only when it supports a safe and achievable transport task.

The safety context follows the whole chain

The NHVR explains that the Heavy Vehicle National Law applies in participating jurisdictions including Victoria to heavy vehicles over 4.5 tonnes gross vehicle mass. It has not commenced in Western Australia or the Northern Territory, although cross-border application can matter. This paper should not be read as a single rulebook for every vehicle or every state. [HVNL and regulations](#).

NHVR guidance on time-slot bookings identifies risks when delivery requirements pressure drivers to speed or drive fatigued. It also discusses congestion and rushed loading or unloading. The guidance encourages communication when delays arise rather than preserving an unrealistic slot at all costs.

[Time-slot booking risks](#).

Its fatigue guidance addresses relevant chain-of-responsibility parties and drivers, including scheduling and rosters that consider fatigue risk. These are regulatory guidance resources explaining obligations, not an AI scheduling approval. [Fitness to drive: fatigue](#).

Faster coordination must not introduce pressure or imply a verified safe schedule. The actual activity and current requirements still need competent review.

A fictional workflow: prepare an exception brief

A fictional Melbourne distributor has a shipment with an unconfirmed unloading slot. Customer service has received a proposed time; the warehouse record still shows the previous booking. The coordinator prepares a brief for the authorised dispatcher and account owner.

The proposed assistant reads a restricted set of shipment-status and booking records. It lists shipment references, the two conflicting notices, their timestamps and the missing confirmation. It drafts an internal question for the coordinator to review. It does not claim the new slot is accepted or that a driver can safely reach it.

The coordinator checks the sources, confirms the booking through the normal process and asks the dispatcher to resolve transport implications. The account owner approves any customer commitment. During the pilot, all communications remain drafts and no system booking changes automatically.

Use synthetic shipment and contact references. This fictional example shows visible uncertainty: absent confirmation becomes a question, not a guessed event.

What would make this worth continuing?

The following is a decision aid for the fictional trial, not an observed result. Agree the limits with the people responsible before testing.

- **Accepted result:** A coordinator-verified exception brief separating confirmed events, conflicting notices and decisions still needed.
- **Who judges it:** The authorised dispatcher and account owner receiving the brief.
- **Hold and return to the manual route when:** Shipment identity is wrong, confirmation is invented or the draft creates an unsafe instruction or commitment.
- **Evidence for another step:** Compare total gathering, checking and follow-up effort for similar exceptions; keep external waiting time separate.

For a fictional slot change, show the new proposed time and the warehouse's earlier confirmed booking as separate dated entries. Hold the draft's conclusion until the coordinator resolves them through the normal route. The assistant cannot infer that the driver can safely make either slot.

Use the [one-workflow worksheet](#) to record the agreed question, evidence and next decision.

Five states, with evidence at each handover

Understand. Follow an exception from the first notice to verified resolution. Identify who knows the facts, who can change a booking and who owns customer commitments. Record chasing, review and waiting separately. The benefit hypothesis is less duplicated information gathering. The gate is an agreed exception type, authorised data, named owner and baseline. A missing ownership decision must be resolved before automation.

Try. Prepare an offline brief from synthetic or approved records. Include conflicting timestamps, a missing confirmation and two similar shipment references. The benefit hypothesis is faster preparation of the coordinator's picture. The gate is source-backed facts, correct identity matching, explicit uncertainty and manageable checking. A fabricated arrival or commitment is a stop condition.

Repeat. Have another coordinator handle several examples using documented instructions. Include a handover between shifts and an unavailable source. The benefit hypothesis is more consistent briefs when the usual person is absent. The gate is repeatable quality, visible unresolved cases and a manual route. Do not report success by excluding difficult exceptions after seeing the results.

Operate. With approval, introduce routine drafting for that exception type. Name an owner and backup, keep access limited and make failures and ageing drafts visible. The benefit hypothesis is less chasing around recurring coordination. The gate is exercised recovery, accepted operating burden and clear limits on sending and system changes. Dispatch authority does not pass to the assistant.

Improve. Compare actual effort, quality and recipient usefulness. A standard booking-confirmation field may remove more friction than additional generation. The benefit hypothesis is retaining changes that improve coordination. The gate is a recorded retain, revise or stop decision and fresh assessment before broader use. Expanding to allocations would require an entirely different scope and authority.

Understand / Try / Repeat / Operate / Improve is an original Alongside People discussion aid, not transport accreditation, a validated score or an official safety framework.

A proposed 30-day pilot

Days 1-7: select one exception type, observe manual handling and obtain data and tool permissions.

Days 8-14: prepare paired briefs offline from matched examples; verify every event and reference. Days

15-23: repeat with another coordinator, introduce a changed booking and exercise a source outage.

Days 24-30: review evidence with dispatch, warehouse, customer service and safety owners.

A pilot can finish offline. The suggested month is not a delivery commitment, promised saving or approval to ingest operational records. Permission or source-quality issues change the sequence.

Measure resolved coordination

Choose active coordinator effort to produce a verified exception brief as the primary measure. Include gathering, prompting, checking, correcting and follow-up. Keep elapsed resolution time separate because it may depend on an external confirmation beyond the team's control.

Guardrails include wrong shipment identity, invented or stale events, missing sources, unapproved commitments, review burden and inappropriate information disclosure. Track whether the recipient still

has to chase the same facts. Include tool costs, integration, training and maintaining records in the operating cost.

Agree thresholds before the trial and preserve rejected examples. Any unsafe instruction, invented confirmation or change outside permission triggers a hold and manual handling. Check whether messages or targets could create pressure even if the draft contains no explicit speeding request. Do not use delivery punctuality as the only success measure.

Recovered capacity may allow coordinators to address a complex exception earlier. It does not establish safer driving, fewer claims, improved on-time performance or cash savings without separate evidence. Avoid transferring the checking burden to dispatch while reporting only the coordinator's gain.

Permissions, recovery and our role

The initial assistant can read the approved record set and create a draft. It cannot contact a driver or customer, alter a booking, assign work, modify a route, approve a load or interpret a work diary. Keep an authoritative status system and retain timestamps and source links.

If a source fails or identity is uncertain, hold the draft and expose the missing fact. Resume the established manual process. Record what was completed so recovery does not duplicate messages or commitments. Test that a backup coordinator can reconstruct the exception from the original records. Correct any circulated error through the authorised operational route.

Our team can help map exception handling, test alternatives and implement the agreed preparation process. We can connect authorised information, make gaps visible, establish review and recovery steps and hand over to the people operating it. Your organisation retains transport, customer and safety authority; wider integration and support are separately scoped.

Official sources inspected 3 October 2026. Recheck current HVNL provisions and jurisdiction before implementation. The sources support the narrow regulatory statements; the workflow and pilot are our proposals, not legal assurance.

If you are already trying AI and are unsure what deserves further investment, [tell us what you have tried and what you want to improve](#). We can help find the starting point together.

Read the structured HTML edition: alongsidepeople.com/resources/ai-for-logistics