

AI & AUTOMATION

AI implementation readiness

A decision guide for the first controlled release

Decide whether a proposed AI workflow is ready for a pilot by testing its value, information boundaries, failure handling and operating ownership.

THE WORK

Define one task, compare it with the current process and test the whole workflow with representative inputs.

THE RESULT

A documented pilot decision that names the permitted use, evidence required, release limits and conditions for stopping.

INSIDE THIS GUIDE

02 Define the task

04 Evaluate the workflow

06 Release within limits

03 Establish the data boundary

05 Design human control

07 Record the pilot decision

01 / PURPOSE

Task definition and ownership

Readiness applies to a specific use of AI. A convincing demonstration cannot answer whether the surrounding workflow is safe or useful.

Write the task as a contract

State the input, expected output, recipient and action that follows. Specify what the system may decide and what remains with a person.

"Summarise approved support notes for an agent" creates a clearer boundary than "automate support".

List the excluded uses. A tool intended to draft an internal response should not quietly acquire permission to send messages, change records or make customer commitments. Changes to those permissions require a new review.

Identify who carries the consequence

Name the people affected by an incorrect, delayed or missing output. Assess whether the action is reversible and how quickly an error can be detected. A draft that an expert checks has a different operating burden from an action that changes a financial or personal record.

Compare with the current process

Observe the work before introducing AI. Measure the time spent on preparation, execution, checking, correction and exceptions. Include waiting time when it affects the customer or a service commitment.

Use the same task definition for the proposed workflow. Count human review, rejected outputs, rework, integration support and recurring operation. Saved drafting time is not the same as available team capacity; the work must actually move or stop.

Set the decision owner

Give one business owner responsibility for the intended outcome and one technical owner responsibility for operation. Involve the relevant data, security and subject-matter owners before building around information or decisions they control.

A readiness blocker

If the team cannot define an acceptable output or explain what happens after a wrong one, narrow the task. A larger model or a better prompt cannot resolve an undefined decision.

02 / DATA

Information boundaries

Trace the information from its original source through retrieval, processing, storage and the final recipient.

01

Source

Identify the owner and permitted purpose.

02

Access

Apply the user's actual permissions.

03

Process

Check vendor handling and retention.

04

Deliver

Control recipients, logs and reuse.

Verify the right to use the inputs

Identify the source, sensitivity, owner and allowed use of each input. Confirm that access and contractual terms permit the proposed processing. Remove information that the task does not need; test whether a less sensitive representation is sufficient.

Review provider retention, training use, access controls, deletion routes and processing locations against the organisation's requirements. A general product claim does not establish the configuration used by your account or integration.

Preserve permissions through retrieval

A search or retrieval layer should not expose a document simply because it is available to the service account. Apply the requesting user's access and test the boundaries with accounts that have different permissions.

Treat retrieved text, uploaded files and external pages as untrusted inputs. They may contain instructions that conflict with the intended task. Keep authority in the application's controls and restrict which tools or actions content can trigger.

Check the hidden copies

Inspect prompts, traces, caches, support exports and monitoring records. They may retain information after the original document has been removed. Give each necessary copy an owner, access rule and retention decision.

03 / EVIDENCE

Workflow evaluation

A useful evaluation covers ordinary work, difficult cases and failures. It tests the surrounding controls as well as the model's output.

01

Build a representative test set

Collect authorised tasks from the intended workflow, with realistic variation in length, language, ambiguity and source quality. Include missing information, conflicting records and requests outside the permitted task. Keep a separate held-out set for the release decision.

Completion check

Every important task type and consequential failure mode has a test, an expected result and an accountable reviewer.

02

Define quality before testing

Choose task-specific criteria: factual support, completeness, correct routing, usable format and permitted action. Define what counts as an unacceptable failure. Reviewers should apply the same rubric and resolve disagreements explicitly.

Completion check

A person can judge an output against the rubric without knowing which model or prompt produced it.

03

Test connected actions

Verify access checks, tool permissions, retries, duplicate prevention, timeouts and record updates. Test attempts to redirect the workflow through hostile content. Confirm that the system stops or routes to a person when required information is unavailable.

Completion check

A failed step cannot silently create a successful status or trigger the same consequential action twice.

04

Compare complete outcomes

Compare the current process and proposed workflow on the same task set. Record quality, completion time, reviewer effort, rework and exception volume. Segment results where a strong average could hide a weak task or user group.

Completion check

The release decision uses measured workflow results and unresolved failures, with the tested configuration recorded.

04 / OVERSIGHT

Human review and authority

Human review works when the reviewer has enough context, time and authority to change the outcome.

Choose where approval belongs

Place a review before the point where a consequential action becomes difficult to reverse. State which outputs always require approval, which may proceed under defined conditions and which must be rejected or escalated.

Avoid using a model's self-reported confidence as the sole permission to act. Tie routing to observable conditions and verified performance for the task. Missing sources, policy conflicts or out-of-scope requests should have explicit paths.

Budget the review capacity

Estimate the number of items arriving, the share requiring review and the time each route consumes. Include rework and specialist escalations. Test the queue under the expected peak workload, not only with a few carefully selected requests.

Provide the evidence to review

Show the original request, relevant source material, proposed output and the action that approval will trigger. Make uncertainty and missing information visible. A polished answer alone can make unsupported statements difficult to detect.

The reviewer must be able to edit, reject or escalate, and to see whether the action actually completed. Record the decision separately from the model's suggestion so later review can distinguish the two.

Define the unavailable path

Set the behaviour when no qualified reviewer is available or a case has waited too long. Use an approved manual route, hold the action or return a clear status. Silence and automatic approval are poor substitutes for an explicit decision.

Prove the stop

The operating owner should be able to disable the AI step, revoke its action permissions and route work through the agreed fallback. Test that route before a pilot includes live users.

Pilot scope and release controls

A pilot should make the remaining uncertainty measurable while keeping its consequences within an agreed limit.

Specify the release envelope

Record the user group, task types, data sources, integrations and actions included in the pilot. Set the operating period and review points. Use access controls to enforce the boundary instead of relying only on an instruction to users.

Document the tested model, prompt, retrieval configuration and tools. A material change can invalidate the evidence used to approve the pilot. Retest the affected behaviour before expanding the release.

Train pilot users on the permitted task, review responsibility, failure reporting and fallback. Check that they can recognise an unsupported result and stop before taking the wrong action.

Define stopping conditions

Set explicit conditions for pausing the workflow: an unauthorised action, access leakage, an unacceptable decision, repeated unhandled failures or loss of review capacity. Choose conditions that fit the task's consequence and operating policy.

Monitor outcomes and workload

Track the same quality and effort measures used in evaluation, together with incidents, overrides, queue age and fallback use. Assign someone to review the results and investigate unexpected changes.

Make failures easy to report and link reports to the relevant run without exposing unnecessary sensitive content. A rising correction rate or a growing review queue may matter before a headline success rate changes.

Decide what expansion requires

Expand only after the owner reviews the actual pilot evidence and remaining gaps. New data, user groups, languages or action permissions need corresponding evaluation. Retain the ability to reduce scope or return to the previous process.

Three defensible decisions

Proceed when the evidence and operating controls support the defined pilot. Hold when a necessary dependency or test remains unresolved. Redesign when the task cannot achieve a useful outcome within an acceptable operating boundary.

06 / DECISION RECORD

Pilot decision record

Complete this record with the business, technical and data owners. Refer to the controlled evaluation evidence; keep sensitive test material in its approved location.

Task and accountable owner

Decision date and reviewer

Permitted users, inputs and actions

Define the scope that the release controls will enforce.

Evidence supporting the decision

Record the evaluation version, key results and unresolved failure modes.

Human review, fallback and stop conditions

Proceed, hold or redesign

Next review and expansion condition

Revisit the decision when the model, data, permissions or intended use changes materially.