



AI automation

Automating repeatable work while preserving human ownership for exceptions and approvals.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 05 / EMBER & TEAL

EXECUTIVE WORKING EDITION - 08 OCT 2026

05

AT A GLANCE

Automating repeatable work while preserving human ownership for exceptions and approvals.



AI automation combines traditional process automation with language and pattern-recognition capabilities so that more of the repeatable work around intake, routing, reporting, reconciliation and follow-up can run without manual effort. The gain is real only when the process is understood, the exceptions are owned and the results are measured. This playbook explains how to choose what to automate, how to separate rules from judgment, how to design exception handling and approvals, and how to prove the return before scaling.

WHO IT IS FOR

Operations, finance, shared-services and technology leaders responsible for processes with high volume, repeated steps and pressure on cycle time or cost.

TIME TO USE

60 minutes to read; one process-mapping session to apply



What you will have at the end

- A scored shortlist of automation candidates, with the reasons for the top pick
- A process map that separates rules, judgment and exceptions
- An exception and approval design with named human owners
- A baseline and a return calculation that includes ongoing support cost
- A 90-day plan from one automated process to a repeatable method

Contents

01	What AI automation is, and what it is not	03
02	Where AI automation fits and where it does not	04
03	How AI automation works, in plain language	06
04	Operating model, roles and ownership	08
05	Risks, controls and human review	09
06	Measuring quality and value	11
07	Implementation guide: map, simplify, automate, measure	13
08	Common pitfalls, an illustrative scenario and vendor questions	14
09	A first quarter that produces evidence	16
10	Terms used in this publication	17



EXECUTIVE BRIEF

What AI automation is, and what it is not

01



Automation moves work from people to systems. AI widens the range of work that can move, mainly by handling language and variation. It does not remove the need to own the outcome.

Process automation has existed for decades: scripts, scheduled jobs, workflow engines and robotic process automation that mimics clicks. These tools excel when inputs are structured and rules are explicit. They struggle when the input is an email, a free-text form, a document in an unexpected layout or a request whose meaning must be interpreted. AI adds a layer that can read, classify, extract, summarize and draft, so a workflow can accept messy input and still produce structured output.

The business case is straightforward but easy to overstate. Savings come from removing handling time, reducing delays and cutting rework, not from the existence of a tool. A process that takes ten minutes of attention per item across a thousand items a month has a clear prize. A process with fifteen items a month does not, however elegant the automation. Equally, automation that moves work from one person to another, for example from clerks to a specialist who now fixes the exceptions, may save little. Count the total effort before and after.

- Use deterministic rules wherever the logic is known; use AI only for the parts that need interpretation
- Design the process around exceptions, because they decide whether automation saves time or creates work
- Keep a named human accountable for the outcome of every automated process
- Measure end-to-end results: cycle time, error rate, rework and cost, not number of automations built

CORE IDEA

Automate the routine path completely, route the unusual path to a skilled person quickly, and make it easy to see which is which.



WORKING QUESTION

Which process in your organization has the highest volume of repeated manual handling, and who owns its outcome?

WHERE IT FITS

Where AI automation fits and where it does not

02

Pick processes by volume, repetition, clarity of rules and cost of error. Shortlist by scoring, not by enthusiasm.

Good candidates share traits: they run frequently, follow a mostly stable pattern, have clear inputs and outputs, and produce a result that can be checked. Poor candidates are rare, depend on tacit judgment, change constantly or carry consequences that need personal accountability. Between these, there are processes that suit partial automation, where AI prepares and a person decides. Often that is the right first step, because it builds trust and produces data on how well the AI performs.

PROCESS	FIT	TYPICAL AUTOMATED HUMAN PART	
Customer or supplier request intake	Strong	Classify, extract details, route to queue	Handle unusual or sensitive requests
Weekly performance reporting	Strong	Gather data, draft commentary	Review narrative and approve
Account or data reconciliation	Strong with controls	Match items, flag differences	Resolve exceptions, approve adjustments
Follow-up reminders and status updates	Strong	Trigger and draft messages	Handle replies needing judgment
Contract or policy exception decisions	Weak	Gather context	Decide and sign

PROCESS	FIT	TYPICAL AUTOMATED HUMAN PART PART
Rare, high-value negotiations	Not suitable	Background research Everything else only

A simple scoring approach

FACTOR	SCORE HIGH WHEN	SCORE LOW WHEN
Volume and frequency	Hundreds of items per month or more	A few per month
Standardization	Same steps in most cases	Every case differs
Data readiness	Inputs are digital and accessible	Paper, email chains, locked systems
Error tolerance	Mistakes are caught later and cheap to fix	Mistakes are costly or irreversible
Process stability	Rules rarely change	Rules change monthly

COMMON MISTAKE

Automating a broken process. If steps exist only to correct earlier errors, remove the cause first. Automation of waste is faster waste.

WORKING QUESTION

If you scored your top five candidate processes against these factors, which would lead, and would that surprise your team?

HOW IT WORKS

How AI automation works, in plain language

03

A typical automated process is a pipeline with triggers, AI steps, rule steps, checks, human handoffs and a record of what happened.

Most successful implementations are not a single clever model but a sequence of small, testable steps. A trigger starts the process: an email arrives, a form is submitted, a file lands in a folder, a schedule fires. An AI step interprets the input, for example classifying the request type or extracting amounts and dates. Rule steps apply known logic, such as approval limits or routing tables. Validation steps check the AI's output against reference data. Where confidence is low or a rule fails, the item goes to a person with context. Finally, the system updates records and logs the result.

01 Trigger

Something starts the run, such as a new message, a file or a scheduled time.

02 Interpret

An AI step classifies, extracts or summarizes the input into structured fields.

03 Validate

Code checks the fields: formats, ranges, known customers, totals that add up, duplicate detection.

04 Decide by rule

Business rules route the item: proceed, request more information, approve automatically within limits, or escalate.

05 Act and record

The system updates records, sends messages and writes a log entry that links back to the original input.

06 Hand off exceptions

A person receives the case with the AI's reading, the checks that failed and a suggested next step.

TOOL TYPE	STRENGTH	LIMIT
Workflow or integration platform	Connects systems, manages triggers and approvals	Needs clear rules; not good at interpreting text
Robotic process automation	Works with systems that lack interfaces	Fragile when screens change
Language model step	Interprets and drafts from messy input	Variable output; needs validation
Classical machine learning	Prediction on structured data at scale	Needs history and monitoring
Scripts and scheduled jobs	Simple, cheap and predictable	Maintenance depends on whoever wrote them

DESIGN PRINCIPLE

Wrap every AI step in validation. Treat its output as a proposal that code and, where needed, a person must confirm before it changes a record.

WORKING QUESTION

In your best candidate process, which single step needs interpretation, and which steps are really just rules?

OPERATING MODEL

Operating model, roles and ownership

04

Automations are living services. They need an owner who answers for results, a maintainer who keeps them working and a clear route for exceptions.

Many organizations build automations quickly and then discover that nobody maintains them. A form changes, a supplier alters an email format or a policy limit moves, and the automation silently misroutes items for weeks. Treat each automated process as a product with a lifecycle. Register it in an inventory with its owner, purpose, systems touched, data used, approval limits and review date. Retire it when it stops paying for itself.

ROLE	OWNS	QUESTIONS THEY ANSWER
Process owner	Business outcome and exceptions	Is it working? Who handles what it cannot?
Automation lead	Method, standards, inventory	Is it built and tested to our standard?
Builder or integrator	Configuration and code	Why did this run fail?
Exception handler	Cases routed to humans	What is the right decision for this item?
Control owner (finance, risk, audit)	Approval limits and segregation of duties	Do controls still hold when the process runs unattended?
Security lead	Credentials, access, data	Does the automation have only the access it needs?

What belongs in an automation inventory

- Name, purpose and the business outcome it supports
- Owner, backup owner and the team handling exceptions
- Systems and data touched, with the access used and the credential owner
- Approval limits and controls built in, such as maximum amounts and duplicate checks
- Volume, cycle time, exception rate and cost per item, updated monthly

- Last test date and next review date

Give exception handlers real authority and good tooling. If handlers must re-key information or hunt for context, the time saved on the main path will be lost on the edge cases. A good exception record shows the original input, what the system extracted, which checks failed and what action is suggested. Track how long exceptions wait, because a growing queue is the first sign that the automation is drifting from reality.

WORKING QUESTION

If the person who built your most important automation left tomorrow, who would notice a failure and who could fix it?

RISKS AND CONTROLS

Risks, controls and human review

05

Automation scales mistakes as efficiently as it scales success. Controls should limit the size of any error and make it quick to detect.

Risk grows with autonomy and with the value of the action. Fraud, privacy and financial reporting rules may constrain what can run unattended, and frameworks such as the NIST AI Risk Management Framework and ISO/IEC 42001, together with data protection law and the EU AI Act where applicable, shape governance expectations. This playbook is not legal advice. Confirm current requirements for your sector, involve your control owners early and document why each approval limit was chosen.

RISK	EXAMPLE	CONTROL
Silent failure	A changed input format causes misrouting for weeks	Volume and exception-rate alerts; daily health checks
Wrong extraction	Amount misread and posted	Validation against totals and reference data; review above threshold

RISK	EXAMPLE	CONTROL
Duplicate or runaway action	Same message sent repeatedly	Idempotency checks; rate limits; kill switch
Segregation of duties erosion	One automation both creates and approves	Separate identities and approvals; control review
Excess access	Service account with broad permissions	Least privilege; credential rotation; access reviews
Over-trust	Staff stop reviewing exceptions carefully	Sampling of automated items; seeded test cases

Control checklist before release

- Every action has a defined maximum value or scope
- Automated decisions above that limit go to a named approver
- A complete log links each outcome to the original input
- Alerts exist for volume spikes, drops and rising exception rates
- A tested way to pause the process and reverse recent actions
- Credentials are dedicated, minimal and owned
- Sample audits are scheduled, with results reported to the process owner

EXAMPLE (ILLUSTRATIVE)

An accounts team automates supplier statement matching. Items matching exactly are cleared automatically up to a set amount. Partial matches go to a clerk with the AI's suggested pairing. A weekly sample of auto-cleared items is checked by a supervisor, and any error pauses the rule.

WORKING QUESTION

What is the largest single error your automation could make before anyone noticed, and how would you cap it?

MEASURING VALUE

Measuring quality and value

06

Measure the whole process before and after, including exceptions, support effort and the cost of errors.

Return on automation is easy to inflate. Count the time people spent on the routine path, and then count the time they now spend on exceptions, monitoring, maintenance, rework and training. Include licence and platform costs, build effort and the cost of errors that reach customers or ledgers. Use observed data rather than estimates where you can: sample real items, time them and record the outcome. Do the same afterward, using the same method, so that the comparison is fair.

METRIC	DEFINITION	WHY IT MATTERS
Straight-through rate	Share of items completed with noMain driver of savings human touch	
Exception rate	Share routed to people	Predictor of hidden work
Accuracy of automated outcomes	Share correct in audit sample	Quality of the routine path
Cycle time	Time from trigger to completion	Customer and internal experience
Cost per item	Total running cost divided by items processed	Unit economics
Rework rate	Items needing correction after completion	Hidden cost
Exception ageing	Time exceptions wait before action	Early warning of overload

Also measure what automation does to the people and customers on either side of it. If a supplier now receives faster confirmations, or an internal requester stops chasing status, that is value, though it is harder to price. Capture it with a short survey or a count of status-chasing messages before and after. Be equally alert to harms: a customer who receives a wrong automated message costs far more goodwill than the minutes saved, so include complaints and corrections in your scorecard from the first week.

Return calculation, step by step

01 Baseline

Measure volume, handling time per item, error rate and rework for a representative period.

02 Project

Estimate straight-through rate from a pilot on real items, not from a demonstration.

03 Add costs

Include build, platform, monitoring, maintenance, exception handling and training.

04 Compare

Calculate net time and cost saved per month, and the payback period.

05 Re-measure

After launch, repeat the measurement at 30, 60 and 90 days and correct the model.

ILLUSTRATIVE FIGURES ONLY

Suppose 2,000 items a month take 6 minutes each, which is 200 hours. If 70 percent run straight through and exceptions take 10 minutes, handling falls to 600 items times 10 minutes, or 100 hours, plus 15 hours of monitoring. The net saving is 85 hours. Use your own data.

WORKING QUESTION

What does your automation case look like once exception handling and maintenance are counted honestly?

IMPLEMENTATION GUIDE

Implementation guide: map, simplify, automate, measure

07



Do the unglamorous process work first. The technology is the shorter part of the project.

01 Map the real process

Walk through ten to twenty actual cases with the people who do the work. Record steps, systems, delays, exceptions and workarounds.



02 Simplify

Remove steps that exist only for historical reasons. Fix upstream causes of errors where possible.

03 Separate rules from judgment

Mark each step as rule, interpretation or decision. Rules become code, interpretation may use AI, decisions stay with people.

04 Design exceptions and approvals

Define limits, thresholds, queues, owners and service levels for human handling.



05 Build, test and shadow

Run on historical items, then in shadow alongside people, and compare outcomes.

06 Release in stages

Start with a share of volume or a low-value band, expand as accuracy holds.

07 Operate

Monitor, review exceptions weekly, update rules and re-test after any change in systems or policy.

Pilot readiness

- Process owner and exception handlers are named and available
- Baseline measured on a representative sample
- Rules, thresholds and approval limits are documented and agreed with control owners
- Test set includes ordinary, unusual and deliberately difficult items



- Monitoring and alerts are in place before go-live
- Rollback and pause procedures are tested
- Handlers have been trained on the new queue and tools



Think about people throughout. Staff who currently do the routine work will be asked to take on exceptions, checks and improvement. That is often more interesting work, but it requires different skills and a different pace. Explain what will change, involve the team in mapping and testing, and plan training for the new roles. Automation programs that ignore this tend to meet quiet resistance, workarounds and poor exception handling, which undermines the results.

WORKING QUESTION

Which steps in your chosen process exist only because of an old system or a past error, and could be removed before automating?



PITFALLS AND VENDORS

Common pitfalls, an illustrative scenario and vendor questions

08



Most automation disappointments trace to weak process understanding, ignored exceptions and no ownership after launch.

- Automating before understanding the process, then discovering a dozen undocumented exceptions
- Counting only the saved time on the happy path
- Giving automations broad system access for convenience
- Skipping monitoring because the early results looked good
- Building many small automations nobody inventories or maintains
- Treating AI output as final when a simple validation would catch most errors



A useful discipline is a quarterly automation review. For each process in the inventory, ask four questions: is the volume still what we assumed, has the exception rate moved, has anything upstream or downstream changed, and would we build this again today? Retire or redesign

what fails the test. Organizations that skip this review accumulate fragile automations that nobody dares to touch, and the maintenance burden slowly consumes the savings the program was meant to deliver.

ILLUSTRATIVE SCENARIO (HYPOTHETICAL)

A distributor automates order intake from emailed purchase orders. In testing it handles most orders. After launch, a customer changes its template and quantities are read from the wrong column. For two weeks, orders ship wrong amounts. A check comparing extracted totals with stated totals, plus an alert on rising exception rates, would have caught it on day one. The distributor adds both.

VENDOR QUESTION	WHAT TO LOOK FOR
Can we run our own sample data through before contracting?	Yes, with transparent results and failure examples
How are credentials and permissions managed?	Dedicated service accounts, least privilege, audit logs
What monitoring and alerting is included?	Run history, error handling, volume alerts
How do you handle exceptions and human tasks?	Built-in queues with context and service levels
How does pricing scale with volume and runs?	Predictable model; no surprise charges
How do updates affect our automations?	Change notice, version control, ability to test first

Platform choice

- Connects to the systems you actually use today
- Supports both rule-based and AI steps with validation
- Provides logs your auditors can read
- Allows your staff to maintain simple changes safely
- Has an exit path for your workflows and data

WORKING QUESTION

What would be the first sign that one of your automations had silently started to fail?

30 / 60 / 90

A first quarter that produces evidence

09

Automate one well-understood process end to end, measure honestly, then turn the method into a repeatable practice.

First 30 days: ground it

- Score candidate processes and choose one with a named owner
- Map twenty real cases, including exceptions and workarounds
- Measure baseline volume, handling time, error and rework
- Agree approval limits and controls with control owners

Days 31 to 60: test it

- Build the pipeline with validation after every AI step
- Test on historical items, then run in shadow mode
- Design exception queues, owners and service levels
- Set up monitoring, alerts and a pause procedure

Days 61 to 90: decide

- Release in stages, starting with low-value items
- Re-measure at 30, 60 and 90 days against the baseline
- Start the automation inventory and review calendar
- Decide scale, change or stop, and pick the next candidate using the same method

GLOSSARY

Terms used in this publication

10

TERM	MEANING
Straight-through processing	Completing an item with no human intervention.
Exception	An item that does not fit the routine path and needs human handling.
Robotic process automation	Software that mimics user actions in existing applications.
Workflow engine	A tool that manages steps, rules, approvals and handoffs.
Validation	Automated checks that confirm extracted or generated data is plausible and consistent.
Approval limit	The maximum value or scope an automation may act on without a person.
Segregation of duties	Separating tasks so no single actor can create and approve the same action.
Idempotency	Designing an action so repeating it does not cause duplicate effects.
Shadow mode	Running a process alongside people without acting, to compare results.
Cycle time	Elapsed time from start to completion of a process.
Rework	Effort spent correcting an item after it was marked complete.
Automation inventory	A register of automated processes with owners, controls and review dates.

A note on use

This publication is a public orientation resource, not legal, technical, financial or professional advice. Adapt it to your people, information, workflows and obligations, and check current tool terms, regulations and risks before acting.



TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.

