

The 15-Minute AI Automation Audit

Choose one workflow worth testing.
Know what to leave alone.

Turn recurring work into a shortlist of safe, useful improvements. Leave with a first pilot, a rough value estimate and a 30-day plan - before buying another tool.

Not sure what to write? Start with the completed guide

A fully completed, fictional example is on pages 9-15, after a blank separator page. Copy its wording pattern, then use your own details and numbers. Start with one task you do most weeks. Rough notes are enough.

The 15-minute route

Time	What to do	Your output
1 min	Set the business goal on this page	One measurable result
3 min	List up to five workflows on page 2	A focused inventory
2 min	Apply the stop gate on page 3	Hold unsafe candidates
3 min	Score eligible workflows on page 4	Value and readiness scores
3 min	Use page 5 to rank and estimate the leader	One provisional pilot
2 min	Shortlist on page 6	Up to three opportunities
1 min	Commit to the first step on page 7	An owner and a date

Business / date

Audit owner

One result I want to improve (e.g. fewer missed enquiries or 2 hours freed weekly)

Type into the fields or print and write. Calculate scores and estimates manually. Mark unknowns; validate with real examples during the 30-day pilot. This is a provisional shortlist, not approval to go live.

Inventory the work

Choose up to five recurring workflows. Write a specific trigger and a finished result: "new enquiry arrives to reply approved" is useful; "marketing" is too broad. Count active work time, not waiting time.

ID	Workflow trigger to finished result	Owner + current tool	Runs / week	Minutes / run	Pain or failure
A					
B					
C					
D					
E					

Starter ideas: enquiry replies, appointment reminders, quote follow-ups, order updates, stock alerts, staff onboarding or weekly reporting. Choose the ones you actually repeat.

Check the workflow before choosing the tool

Remove then simplify then automate

Can you stop doing it? Can a template, clearer handover or existing software feature solve it? Use ordinary automation for fixed rules; consider AI where language or varied inputs need interpretation.

Leader to inspect first (ID)

Weekly active hours = runs x minutes / 60

Example only: reminders: 30 runs x 10 minutes / 60 = 5 hours/week. Monthly work? Divide monthly runs by 4.3 to estimate weekly runs. Use the same start-to-finish scope and avoid counting shared steps twice.

The do not automate gate

For every candidate, ask each question below. Any YES or UNKNOWN means HOLD: do not connect live data, auto-send messages or let the system act. A high score on page 4 cannot override this gate.

Red flag	Ask this before any live pilot
Critical judgement	Would this replace a qualified person or accountable human in health, safety, legal advice, hiring, credit or another high-consequence decision?
Uncontrolled authority	Could it spend or move money, change bank details, grant access, delete records, or make a binding promise without approval and limits?
Unapproved data	Would it expose personal, confidential or sensitive information to a tool without permission, approved access and an agreed storage/deletion approach?
No reliable evidence	Are source data or instructions unreliable, or is there no practical way for a person to check the result before harm occurs?
Hidden exceptions	Are important exceptions unclear, or could an error cause serious harm before anyone notices it?
No owner or recovery	Is there no named person to review errors, stop the workflow and restore the manual process?

HOLD workflow IDs + reason

Owner + what must change before reconsidering

Safe next step for a held workflow

Map or redesign the process using made-up or properly de-identified examples. Restore human authority, resolve permission gaps and document checks. Re-run the gate before any live use. If critical judgement remains, keep it human-led.

Treat emails, files and web pages as untrusted input. An AI suggestion must not grant itself permission to take an action. Restrict access and approve consequential actions separately.

Score opportunity and exposure

For workflows that passed the gate, choose 0, 1, 2 or 3 for each dimension. These are practical ranking bands, not a validated prediction. Never add risk or data sensitivity to the opportunity score.

Dimension	0	1	2	3
T Time / week	Under 0.5 hour	0.5 to under 2 hours	2 to under 5 hours	5+ hours
P Repetition	Less than weekly	1-4 runs / week	5-19 runs / week	20+ runs / week
L Rules	No stable process	Main steps known; exceptions unclear	Written steps; exceptions need judgement	Clear rules; exceptions routed
F Automation fit	No usable input or checkable output	Inputs vary or tool fit unknown	Usable digital inputs; tool fit plausible	Stable inputs; small test proves fit
R Error risk	Internal and trivial to fix	Reversible; limited disruption	Customer, money or important records affected	Serious harm or critical decision
S Data sensitivity	Public or made-up data	Routine business data or contact details	Confidential or personal records	Highly sensitive data or secrets

Opportunity value = T + P (0-6). Build readiness = L + F (0-6). Risk and sensitivity are separate controls. Unknown T/P/L/F? Use 0 provisionally and verify; unknown R/S? Return to HOLD.

ID	T	P	L	F	R	S	Value T+P	Ready L+F
A								
B								
C								
D								
E								

Control rules: R or S = 3: HOLD for qualified review and a redesigned scope. R = 2: a named person checks every output before action. S = 2: document permission, approved tool settings and minimum data access. Lower scores still need checks; the gate always applies.

Prioritise and estimate impact

High value or readiness = 4-6; low = 0-3. Pilot-first also requires L and F both at least 2; otherwise prepare first. Prefer lower risk and sensitivity, then less setup work and stronger evidence. Estimate the leader today.

Low readiness 0-3		High readiness 4-6
High value 4-6	PREPARE FIRST Clarify rules, clean inputs and route exceptions.	PILOT FIRST Choose one small, reversible workflow.
Low value 0-3	PARK OR REMOVE Stop, simplify or leave it manual.	QUICK WIN IF EASY Use a template or existing feature; cap effort.

Pilot-first IDs

Prepare-first IDs

Quick-win / park IDs

Conservative weekly estimate

Use one currency. B = baseline staff hours/week; f = fraction removed before new oversight (40% = 0.40); O = added review, exceptions and maintenance hours/week. Net hours $H = B \times f - O$. Count all staff; O includes only added overhead.

B Baseline hours	f Removed fraction	O Added hours	H Net hours

r Value per hour	C Weekly added cost	K One-off setup cost	V Weekly net value

$V = H \times r - C$. Use a realistic staff cost or capacity value for r. K includes setup labour; C includes extra tool/usage costs. Value payback = K / V weeks when V is positive. If H or V is not positive, redesign or park.

Worked example - illustrative, not a promised result
Reminders: B = 5 hours, f = 0.40, O = 0.5 hour. H = 1.5 hours/week. At \$50/hour and C = \$10/week, V = \$65/week. K = \$300 gives about 4.6 weeks value payback. If f is only 0.20, V is \$15 and payback becomes 20 weeks.

Reality check: capacity value is not cash saved. Cash payback requires actual cost reduction or measured extra margin. Do not count extra margin again as time value. Annual net value = $V \times \text{working weeks} - K$ in year one; use your actual working weeks.

Your top three opportunities

Keep only gate-cleared candidates here. Rank the leader first; the others stay in the queue. Fewer than three good candidates is a valid result. Use rough notes now and complete test details before live use.

Opportunity 1 | first pilot

ID + workflow / finished result		Route: template / rules / AI assist / human
Value / readiness; R / S	Owner + reviewer	One small test + success measure
Estimate: net hours / weekly value	Main uncertainty + how to check it	

Opportunity 2 | queued

ID + workflow / finished result		Route: template / rules / AI assist / human
Value / readiness; R / S	Owner + reviewer	One small test + success measure
Estimate: net hours / weekly value	Main uncertainty + how to check it	

Opportunity 3 | queued

ID + workflow / finished result		Route: template / rules / AI assist / human
Value / readiness; R / S	Owner + reviewer	One small test + success measure
Estimate: net hours / weekly value	Main uncertainty + how to check it	

Example: reminders | scheduling tool | Value 6 / Ready 5 (L3, F2); R2 / S1 | owner checks each recipient and timing before sending | test 10 reminders: no wrong recipients or missed sends; at least 1 net hour saved/week. AI may be unnecessary.

Focus rule: run one pilot at a time. Keep the next two queued until the first has measured benefit, acceptable quality and a reliable manual fallback.

Your 30-day action plan

A pilot succeeds when it saves net time and preserves quality. Assign one accountable owner, a reviewer and a manual fallback. Use this plan after the 15-minute shortlist.

When	Action	Evidence to keep
Days 1-7	Observe at least 10 real runs or every run if fewer occur. Record active minutes, mistakes and exceptions. Simplify first.	A baseline; written steps; agreed acceptance criteria.
Days 8-14	Use made-up or approved de-identified data. Test normal cases plus missing data, duplicates, bad input and a tool outage.	Test results; correct outputs; stop and recovery steps.
Days 15-21	Only after the gate passes: pilot a limited live scope. Have a person check every output before any action. Log added review time and costs.	A run log of time, errors, exceptions and costs.
Days 22-30	Recalculate H and V for the same scope. Expand only if targets are met, both are positive and no stop rule is unresolved. Otherwise revise or stop.	A documented decision and the next review date.

Write the acceptance rules before testing

Minimum net hours saved / week

Maximum tolerable error or rework rate

Reviewer + who can stop it

Manual fallback + where the run log lives

Stop rules are mandatory

Any data exposure, unauthorised action, critical error or uncheckable result: stop immediately and use the manual process. A missed target means fix and retest. Do not expand on the strength of one good run.

First action + owner + due date

Day-30 decision: stop / revise / expand

Next review date

WorkPilot Lab standard: keep what is measured to work. Review after any tool, data or process change. This workbook supports an initial business decision; it does not certify security, compliance or professional judgement.

The 15-Minute AI Automation Audit

Choose one workflow worth testing.
Know what to leave alone.

Turn recurring work into a shortlist of safe, useful improvements. Leave with a first pilot, a rough value estimate and a 30-day plan - before buying another tool.

Not sure what to write? Start with the completed guide

A fully completed, fictional example is on pages 9-15, after a blank separator page. Copy its wording pattern, then use your own details and numbers. Start with one task you do most weeks. Rough notes are enough.

The 15-minute route

Time	What to do	Your output
1 min	Set the business goal on this page	One measurable result
3 min	List up to five workflows on page 2	A focused inventory
2 min	Apply the stop gate on page 3	Hold unsafe candidates
3 min	Score eligible workflows on page 4	Value and readiness scores
3 min	Use page 5 to rank and estimate the leader	One provisional pilot
2 min	Shortlist on page 6	Up to three opportunities
1 min	Commit to the first step on page 7	An owner and a date

Business / date

BrightDay Home Cleaning (fictional) / 8 Oct 2026

Audit owner

Sam, business owner

One result I want to improve (e.g. fewer missed enquiries or 2 hours freed weekly)

Free at least 1 staff hour each week on booking reminders, with no missed or wrong reminders.

Type into the fields or print and write. Calculate scores and estimates manually. Mark unknowns; validate with real examples during the 30-day pilot. This is a provisional shortlist, not approval to go live.

Inventory the work

Choose up to five recurring workflows. Write a specific trigger and a finished result: "new enquiry arrives to reply approved" is useful; "marketing" is too broad. Count active work time, not waiting time.

ID	Workflow trigger to finished result	Owner + current tool	Runs / week	Minutes / run	Pain or failure
A	Tomorrow's booking appears in calendar; reminder checked and sent.	Jo / booking calendar	30	10	Typing each message; timing mistakes.
B	Quote is 3 days old; follow-up draft checked and sent.	Sam / email + quote list	15	12	Follow-ups slip when we are busy.
C	New booking confirmed; welcome instructions checked and sent.	Jo / email + saved templates	10	12	Retyping the same instructions.
D	Complaint arrives; owner decides reply and any refund.	Sam / inbox + payment app	2	25	Refund limits and exceptions unclear.
E	Weekly blank checklist opened; formatting changed again.	Sam / document template	1	20	Changing the layout adds no value.

Starter ideas: enquiry replies, appointment reminders, quote follow-ups, order updates, stock alerts, staff onboarding or weekly reporting. Choose the ones you actually repeat.

Check the workflow before choosing the tool

Remove then simplify then automate

Can you stop doing it? Can a template, clearer handover or existing software feature solve it? Use ordinary automation for fixed rules; consider AI where language or varied inputs need interpretation.

Leader to inspect first (ID)

A - booking reminders

Weekly active hours = runs x minutes / 60

30 x 10 / 60 = 5 hours/week

Example only: reminders: 30 runs x 10 minutes / 60 = 5 hours/week. Monthly work? Divide monthly runs by 4.3 to estimate weekly runs. Use the same start-to-finish scope and avoid counting shared steps twice.

The do not automate gate

For every candidate, ask each question below. Any YES or UNKNOWN means HOLD: do not connect live data, auto-send messages or let the system act. A high score on page 4 cannot override this gate.

Red flag	Ask this before any live pilot
Critical judgement	Would this replace a qualified person or accountable human in health, safety, legal advice, hiring, credit or another high-consequence decision?
Uncontrolled authority	Could it spend or move money, change bank details, grant access, delete records, or make a binding promise without approval and limits?
Unapproved data	Would it expose personal, confidential or sensitive information to a tool without permission, approved access and an agreed storage/deletion approach?
No reliable evidence	Are source data or instructions unreliable, or is there no practical way for a person to check the result before harm occurs?
Hidden exceptions	Are important exceptions unclear, or could an error cause serious harm before anyone notices it?
No owner or recovery	Is there no named person to review errors, stop the workflow and restore the manual process?

HOLD workflow IDs + reason

D - refund decisions: spending limits and complaint exceptions are unclear. A, B, C and E: all six answers NO within their stated scope.

Owner + what must change before reconsidering

Sam: keep D manual until refund limits, approvals and exceptions are written. A/B/C use approved contact data only; Jo checks every message before sending.

Safe next step for a held workflow

Map or redesign the process using made-up or properly de-identified examples. Restore human authority, resolve permission gaps and document checks. Re-run the gate before any live use. If critical judgement remains, keep it human-led.

Treat emails, files and web pages as untrusted input. An AI suggestion must not grant itself permission to take an action. Restrict access and approve consequential actions separately.

Score opportunity and exposure

For workflows that passed the gate, choose 0, 1, 2 or 3 for each dimension. These are practical ranking bands, not a validated prediction. Never add risk or data sensitivity to the opportunity score.

Dimension	0	1	2	3
T Time / week	Under 0.5 hour	0.5 to under 2 hours	2 to under 5 hours	5+ hours
P Repetition	Less than weekly	1-4 runs / week	5-19 runs / week	20+ runs / week
L Rules	No stable process	Main steps known; exceptions unclear	Written steps; exceptions need judgement	Clear rules; exceptions routed
F Automation fit	No usable input or checkable output	Inputs vary or tool fit unknown	Usable digital inputs; tool fit plausible	Stable inputs; small test proves fit
R Error risk	Internal and trivial to fix	Reversible; limited disruption	Customer, money or important records affected	Serious harm or critical decision
S Data sensitivity	Public or made-up data	Routine business data or contact details	Confidential or personal records	Highly sensitive data or secrets

Opportunity value = T + P (0-6). Build readiness = L + F (0-6). Risk and sensitivity are separate controls. Unknown T/P/L/F? Use 0 provisionally and verify; unknown R/S? Return to HOLD.

ID	T	P	L	F	R	S	Value T+P	Ready L+F
A	3	3	3	2	2	1	6	5
B	2	2	2	2	2	1	4	4
C	2	2	3	2	2	1	4	5
D	-	-	-	-	-	-	HOLD	HOLD
E	0	1	1	1	0	0	1	2

Control rules: R or S = 3: HOLD for qualified review and a redesigned scope. R = 2: a named person checks every output before action. S = 2: document permission, approved tool settings and minimum data access. Lower scores still need checks; the gate always applies.

Prioritise and estimate impact

High value or readiness = 4-6; low = 0-3. Pilot-first also requires L and F both at least 2; otherwise prepare first. Prefer lower risk and sensitivity, then less setup work and stronger evidence. Estimate the leader today.

Low readiness 0-3		High readiness 4-6
High value 4-6	PREPARE FIRST Clarify rules, clean inputs and route exceptions.	PILOT FIRST Choose one small, reversible workflow.
Low value 0-3	PARK OR REMOVE Stop, simplify or leave it manual.	QUICK WIN IF EASY Use a template or existing feature; cap effort.

Pilot-first IDs

A first; B/C queued

Prepare-first IDs

None; D stays on HOLD

Quick-win / park IDs

Park E; stop reformatting

Conservative weekly estimate

Use one currency. B = baseline staff hours/week; f = fraction removed before new oversight (40% = 0.40); O = added review, exceptions and maintenance hours/week. Net hours H = B x f - O. Count all staff; O includes only added overhead.

B Baseline hours	f Removed fraction	O Added hours	H Net hours
5 hours/week	0.40 (40%)	0.50 hour/week	1.50 hours/week

r Value per hour	C Weekly added cost	K One-off setup cost	V Weekly net value
AUD \$50/hour	AUD \$10/week	AUD \$300	AUD \$65/week

$V = H \times r - C$. Use a realistic staff cost or capacity value for r. K includes setup labour; C includes extra tool/usage costs. Value payback = K / V weeks when V is positive. If H or V is not positive, redesign or park.

Worked example - illustrative, not a promised result

Reminders: B = 5 hours, f = 0.40, O = 0.5 hour. H = 1.5 hours/week. At \$50/hour and C = \$10/week, V = \$65/week. K = \$300 gives about 4.6 weeks value payback. If f is only 0.20, V is \$15 and payback becomes 20 weeks.

Reality check: capacity value is not cash saved. Cash payback requires actual cost reduction or measured extra margin. Do not count extra margin again as time value. Annual net value = V x working weeks - K in year one; use your actual working weeks.

Your top three opportunities

Keep only gate-cleared candidates here. Rank the leader first; the others stay in the queue. Fewer than three good candidates is a valid result. Use rough notes now and complete test details before live use.

Opportunity 1 | first pilot

ID + workflow / finished result		Route: template / rules / AI assist / human	
A - booking reminders checked and sent		Rules: booking reminder template	
Value / readiness; R / S	Owner + reviewer	One small test + success measure	
6 / 5; R2 / S1	Sam owns; Jo checks	10 reminders: no wrong or missed sends; at least 1 net hour/week.	
Estimate: net hours / weekly value		Main uncertainty + how to check it	
1.5 h / AUD \$65 per week (estimate)		Timing errors? Test cancelled and rescheduled jobs.	

Opportunity 2 | queued

ID + workflow / finished result		Route: template / rules / AI assist / human	
B - quote follow-up drafts checked and sent		AI assist: draft only; human sends	
Value / readiness; R / S	Owner + reviewer	One small test + success measure	
4 / 4; R2 / S1	Sam owns and checks	10 drafts: quoted prices unchanged; at least 0.5 net hour/week.	
Estimate: net hours / weekly value		Main uncertainty + how to check it	
0.8 h / AUD \$35 per week (estimate)		Price changes? Compare each draft with the approved quote.	

Opportunity 3 | queued

ID + workflow / finished result		Route: template / rules / AI assist / human	
C - welcome instructions checked and sent		Rules: choose an approved template	
Value / readiness; R / S	Owner + reviewer	One small test + success measure	
4 / 5; R2 / S1	Sam owns; Jo checks	10 emails: correct template/recipient; at least 0.5 net hour/week.	
Estimate: net hours / weekly value		Main uncertainty + how to check it	
0.7 h / AUD \$35 per week (estimate)		Wrong template? Test first and repeat bookings.	

Example: reminders | scheduling tool | Value 6 / Ready 5 (L3, F2); R2 / S1 | owner checks each recipient and timing before sending | test 10 reminders: no wrong recipients or missed sends; at least 1 net hour saved/week. AI may be unnecessary.

Focus rule: run one pilot at a time. Keep the next two queued until the first has measured benefit, acceptable quality and a reliable manual fallback.

Your 30-day action plan

A pilot succeeds when it saves net time and preserves quality. Assign one accountable owner, a reviewer and a manual fallback. Use this plan after the 15-minute shortlist.

When	Action	Evidence to keep
Days 1-7	Observe at least 10 real runs or every run if fewer occur. Record active minutes, mistakes and exceptions. Simplify first.	A baseline; written steps; agreed acceptance criteria.
Days 8-14	Use made-up or approved de-identified data. Test normal cases plus missing data, duplicates, bad input and a tool outage.	Test results; correct outputs; stop and recovery steps.
Days 15-21	Only after the gate passes: pilot a limited live scope. Have a person check every output before any action. Log added review time and costs.	A run log of time, errors, exceptions and costs.
Days 22-30	Recalculate H and V for the same scope. Expand only if targets are met, both are positive and no stop rule is unresolved. Otherwise revise or stop.	A documented decision and the next review date.

Write the acceptance rules before testing

Minimum net hours saved / week	Maximum tolerable error or rework rate
At least 1 net staff hour per week	Zero wrong, duplicate or missed sends
Reviewer + who can stop it	Manual fallback + where the run log lives
Jo reviews each send; Sam can stop it	Send manually from calendar; log in shared Reminders sheet

Stop rules are mandatory

Any data exposure, unauthorised action, critical error or uncheckable result: stop immediately and use the manual process. A missed target means fix and retest. Do not expand on the strength of one good run.

First action + owner + due date
Sam: time 10 reminders, confirm contact-data approval and write timing rules by 14 Oct 2026.

Day-30 decision: stop / revise / expand	Next review date
REVISE: $5 \times 30\% - 0.7 = 0.8$ hours saved; AUD \$30/week. Below 1-hour goal. Keep checks; retest.	20 Nov 2026 (after day-30 review on 6 Nov)

WorkPilot Lab standard: keep what is measured to work. Review after any tool, data or process change. This workbook supports an initial business decision; it does not certify security, compliance or professional judgement.