Choose what is worth automating.

Fictional demonstration / no client engagement or measured results

Morrow Studio is an invented 18-person B2B service business. Every volume, time estimate, cost and projected benefit in this case is a demonstration assumption. No interviews were conducted and no system was deployed. The purpose is to show how I would evaluate operational improvements and challenge an AI business case.

Fictional brief

Requests arrive in a shared inbox, project updates live in spreadsheets and meeting decisions are copied into task lists. The operations lead wants less chasing and clearer reporting. The assessment covers intake, weekly reporting and meeting follow-up; it excludes customer-facing decisions, payroll and payment approval.

Current workflow

01 / Read unstructured input	02 / Copy details into another tool	03 / Chase missing information	04 / Compile an update	05 / Owner reviews and acts

Proposed workflow

01 / Capture consistent source data	02 / Apply rules and completeness checks	03 / Draft a summary if useful	04 / Owner checks facts and decisions	05 / Publish an approved update

Prioritised opportunities

Priority	Opportunity	Judgment
1	Structured weekly reporting	Fix definitions and ownership; test rules first
2	Request intake assistance	Draft and validate; do not auto-commit scope
3	Meeting action drafts	Pilot only with approved data and consent
Defer	Autonomous client replies	External error cost exceeds the case for speed

Why this order

Reporting is recurring, internally visible and easier to compare with a baseline. Intake has more varied language and exceptions. Meeting notes introduce recording and data-handling questions. The priority is based on readiness and risk as well as potential minutes saved.

A business case that can be challenged

Fictional demonstration / no client engagement or measured results

Assumptions per working week

Task	Volume	Before / after	Eligible share	Time released
Intake	40 requests	8 / 5 min	60%	72 min
Reporting	10 updates	18 / 8 min	80%	80 min
Meeting actions	12 meetings	15 / 7 min	75%	72 min

Calculation

Time released = volume × (baseline handling time - future handling time) × eligible share. Future handling time includes human review. Tasks are defined as separate activities to avoid double-counting. Gross potential is 224 minutes per week. Subtract 45 minutes for ongoing maintenance: net potential is 179 minutes, or about 3.0 hours per week.

Capacity, not cash savings

At an assumed 48 working weeks, the base case releases 143.2 hours a year. This is available capacity, not a payroll reduction or realised financial return. Implementation effort, software costs and how the capacity would be used must be assessed before approval.

Sensitivity before a purchase

Scenario	Gross benefit realised	Net min/week	Net hours/year
Conservative	50% of base estimate	67	53.6
Base assumption	100% of base estimate	179	143.2

What would invalidate the case

The workflow is infrequent, input data is unreliable, review takes as long as the existing task, exceptions dominate, or no owner can maintain it. In the conservative scenario, 45 minutes of maintenance still apply. Actual timings and eligible shares must be observed before relying on either scenario.

A small pilot with a clear stop rule

Fictional demonstration / no client engagement or measured results

Pilot specification

Start with weekly reporting only. Use one agreed source table with project ID, owner, reporting period, status, next milestone and blocker. Check missing fields and duplicate project-period records with rules. AI drafting is optional; a template may be sufficient. Any draft must reference its source rows and label missing information.

Review and failure handling

The operations lead approves every update before publication. Missing owners, conflicting dates and unsupported statements go to an exception queue. No autonomous changes to project commitments or external messages. Keep an editable log and a manual fallback. Use only client-approved data and tools; redact sensitive inputs during testing.

Thirty-day sequence

When	Work	Decision
Days 1-5	Observe a full cycle; agree baseline	Is the problem frequent and measurable?
Days 6-10	Clean definitions; prototype one flow	Can ordinary rules do the job?
Days 11-20	Shadow-run alongside manual process	Are outputs complete and supportable?
Days 21-30	Review effort, errors and ownership	Adopt, revise or stop

Proposed acceptance criteria

Every published statement is supported by source data and approved by its owner. All required fields are present or explicitly flagged. Run at least two reporting cycles and compare handling time including review and corrections. Progress only if net effort falls, no critical factual errors reach publication and a named person accepts maintenance ownership.

Audit output

A process map, ranked opportunity list, assumptions table, benefit sensitivity and pilot brief. These are sufficient to make the next decision without committing to a large build. This is a fictional analysis demonstration; projected benefits are not client results.