

JULY 2026 · 9-MIN READ · FOR OPERATIONS, FINANCE & IT LEADERS

Fix the process first.

How to decide what to automate, and what to leave well alone. A method for choosing candidates, mending them before you mechanise, and proving the benefit afterwards.

4 SELECTION TRAPS

ECRS DISCIPLINE

THE FIVE-M METHOD

CLOUDBLISS.CO.UK
UK MICROSOFT CONSULTANCY

THE MOST EXPENSIVE PROJECT

Automation is an **amplifier**, not a **fix**.

The most expensive automation project isn't the one that never ships. It's the one that ships on schedule, works exactly as specified, and makes a broken process run faster.


Bill Gates made this observation decades ago: automation applied to an efficient operation magnifies the efficiency, and applied to an inefficient one it magnifies the inefficiency. The technology has changed several times since (macros, workflow tools, RPA, and now AI agents) but the observation hasn't aged a day.

Point an amplifier at a good process and it multiplies the value. Point it at a bad one and it multiplies the mess, with fewer humans left in the loop to catch it.

Deciding what to automate is a process question, not a technology question. Here is how we answer it at Cloudbliss, and the four traps we see mid-market businesses fall into before they call us.

Three generations of technology, one recurring cause of death. The platforms work. The selection doesn't.

CLOUDBLISS · ON RPA, WORKFLOW AND AGENTS



WHAT'S INSIDE

The four selection traps

PAGE 03

A structured method for choosing

PAGES 04 – 06

Mend before you mechanise

PAGE 07

Build, measure, and the Five-M Method

PAGES 08 – 09

THE NUMBERS BEAR IT OUT

30^{-50%}

OF INITIAL RPA PROJECTS FAIL

EY's observation across early robotic process automation programmes.

3%

SCALED BEYOND ~50 ROBOTS

Deloitte's global survey: almost nobody got past the pilot plateau.

40%+

OF AGENTIC AI PROJECTS CANCELLED BY END 2027

Gartner's forecast: unclear business value, rising costs, weak controls.

BEFORE THEY CALL US

Four ways companies pick the wrong things to automate.

TRAP 01

Tool-first thinking

The licences arrive before the strategy does. "We've got Power Automate, go and find things to automate." What you get is trivia: file-moving flows, birthday reminders, notification chains. Lots of activity, no outcomes.

TELL · MEASURED IN FLOWS BUILT, NOT HOURS RETURNED

TRAP 02

The loudest-voice backlog

The team that complains most articulately gets automated first. But the squeakiest wheel is rarely the biggest constraint. The process quietly haemorrhaging twenty hours a week in finance loses out to the mildly annoying one with the most persuasive manager.

TELL · PRIORITIES TRACK SENIORITY, NOT VOLUME

TRAP 03

The wishlist workshop

A room, a hundred sticky notes, and a prioritisation exercise by show of hands. Workshops are useful for surfacing candidates. They are useless for choosing between them, because enthusiasm is not data.

TELL · NO BASELINE FOR VOLUME, TIME OR ERROR RATES

TRAP 04

Paving the cowpath

The most dangerous one. A "process" has accreted over years: a workaround for a 2019 system limitation, a double-entry step nobody questions, an approval that exists because of a forgotten incident. Automate it as-is and you encode every one of them into software.

TELL · NOW THE WORKAROUND HAS AN API

THE COMMON THREAD

All four traps share one root: choosing a candidate before anyone has measured it. Selection is where automation programmes are won or lost, and it happens weeks before the first flow is built.

A STRUCTURED METHOD

There is a better way, and it isn't complicated. It's just disciplined.

01 Start with value streams, not org charts

Inventory processes end to end: order-to-cash, procure-to-pay, hire-to-retain, incident-to-resolution. Departments optimise their own patch. The real pain lives in the handoffs between them, which is precisely where nobody is looking.

Evidence, not anecdote

Timestamps from the systems you already own beat any workshop vote.

02 Follow the actual work, not the documented process

Interviews are necessary but they lie, innocently. People describe the official process; observation and data reveal the performed one, complete with variants, shadow spreadsheets and the step everyone does differently.

Where systems allow, pull the timestamps: ticket data, approval logs, email trails. Where they don't, run a simple one-week sampling exercise. The gap between the documented process and the real one is usually where the opportunity is hiding.

Handoffs first

Cross-team gaps are unowned, unmeasured, and usually the largest prize.

One week is enough

A short sampling exercise gives a defensible baseline without a data project.

03 Quantify, and separate touch time from cycle time

This distinction alone will kill half of a typical automation wishlist, and it should. Score each candidate on volume × touch time, error and rework rates, and the cost of delay. See overleaf for why the two clocks matter.

Steps 04 and 05

Screening and shipping follow on page 06, once the numbers exist.

THE GAP YOU ARE LOOKING FOR

THE DOCUMENTED PROCESS

Six steps. One approval. Two days, according to the process map that was signed off in 2021.


THE PERFORMED PROCESS

Nine steps, three regional variants, a shared spreadsheet nobody owns, and eleven days of elapsed time. This is the version you have to automate, so this is the version worth mending.

STEP 03, IN DETAIL

A ten-day cycle with forty minutes of typing is a waiting problem.

Touch time is the minutes of actual human effort. Cycle time is the elapsed days from start to finish. Automate the typing in the example below and you have saved forty minutes of a ten-day wait. The fix is in the handoffs, the batching and the approval queue, not in a flow.



**Automation removes the effort.
Only redesign removes the delay.**

SCORE EVERY CANDIDATE ON THREE THINGS

Volume × touch time for the size of the prize.
Error and rework rates for the hidden cost. **Cost of delay** for what the waiting is worth. Real numbers only.

Ask which clock hurts

If touch time dominates, automate. If cycle time dominates, redesign the queue first: parallel approvals, thresholds, continuous rather than batch processing. Then automate what's left, on a process worth encoding.

STEPS 04 & 05

Screen for suitability, then ship a thin slice.

The profile of a good candidate

- ✓ Rules-based decisions, not judgement calls
- ✓ Structured, digital inputs
- ✓ Stable process on stable systems
- ✓ Meaningful volume
- ✓ Low exception rate
- ✓ A clear, named owner

In 2026, AI agents have genuinely extended this frontier. Document extraction, triage, first-draft generation and other judgement-adjacent work are now automatable in ways they weren't three years ago.

But note what hasn't changed: agents raise the bar on process clarity, they don't lower it. An AI agent pointed at an ambiguous process produces ambiguity at scale, with confidence.

MATCH THE TOOL CLASS TO THE PROBLEM

DETERMINISTIC FLOW

Rules and structured data

Known inputs, known steps, an auditable path. Power Automate territory. Cheapest to run, easiest to trust.

EXAMPLE · SUPPLIER ONBOARDING CHECKS

AI AGENT

Language and documents

Unstructured input, judgement-adjacent tasks, a human check on the output. Copilot Studio and AI Builder territory.

EXAMPLE · TRIAGING AN INBOX OF QUERIES

HUMAN

Ambiguity and consequence

Rare, high-stakes or genuinely contested decisions. Leaving these alone is a design decision, not a failure.

EXAMPLE · A DISPUTED CREDIT DECISION

05 Prioritise ruthlessly and ship a thin slice

Impact against effort, on real numbers. Then pick one candidate that touches genuine volume, deliver it in weeks not quarters, measure it against the baseline, and bank the credibility. Automation programmes live or die on their first three deliveries.

THE STEP ALMOST EVERYONE SKIPS

Mend before you mechanise.

This is the reason we are often hired after a first attempt has disappointed. Fix the process before you build anything, using a discipline borrowed from lean practice: eliminate, combine, rearrange, simplify, applied in that order, because each step is cheaper than the one after it.

CHEAPEST

MOST EFFORT

E ELIMINATE	The best automation is deletion If an approval has been approved 100% of the time for twelve months, it isn't a control. It's a delay with a signature. Remove the reports nobody reads and a surprising share of the backlog dissolves.	Checks that duplicate a system validation. Sign-offs that exist out of habit.
C COMBINE	Merge duplicate capture One form instead of three. One handoff instead of four. Every merge removes a queue as well as a keystroke.	Three departmental forms collecting the same customer reference.
R REARRANGE	Parallel, not sequential Run steps at the same time that currently queue behind each other. Move validation to the point of entry rather than the point of failure.	Two approvers in parallel instead of a relay. Cycle time halves, effort unchanged.
S SIMPLIFY	Standardise before you encode Reduce the variants and the rule set, and fix data quality at the source. Automation doesn't just process bad data faster, it launders it into downstream systems where it is ten times more expensive to correct.	Four regional variants of one approval matrix, reduced to one.

WHY IT MATTERS

Only a simplified, stabilised process gives you a stable requirement. Build on anything else and you'll spend the next year re-platforming your own confusion.

THE STRAIGHTFORWARD PART

Then build, and **measure** like you mean it.

With the process mended, the build is the easy half, and this is where the Microsoft stack earns its keep for mid-market firms.

POWER AUTOMATE

Deterministic flows

The rules-based spine: routing, syncing, posting, chasing.

AUDITABLE · CHEAPEST TO RUN

POWER APPS & FORMS

Clean data at source

Structured capture, so validation happens at entry rather than at failure.

FEWER EXCEPTIONS DOWNSTREAM

COPILOT STUDIO & AI BUILDER

Documents and language

Extraction, triage and first drafts, with a human check where it counts.

NEEDS A CLEAR PROCESS FIRST

Design the exceptions in

An exception route, alerting, a named owner, a runbook. Bots rarely fail loudly. They fail silently, at 2am, in a queue nobody checks.

Right-size the governance

Environments, data loss prevention policies, naming standards, ownership. Enough structure to prevent flow sprawl; none of the centre-of-excellence theatre that suffocates a 200-person business.

MEASURE AGAINST THE BASELINE YOU CAPTURED IN DISCOVERY

MEASURE	AT 30 DAYS	AT 90 DAYS	DECISION
Cycle time	Trend vs baseline	Sustained?	Iterate the handoffs
Error and rework rate	Exceptions logged	Falling?	Fix the rule set
Cost per transaction	First read	Verified	Keep, extend or retire
Adoption	Who is using it	Still using it	Retrain or redesign

Fix or kill what underperforms. Success is counted in hours returned and errors removed, never in the number of flows deployed.

HOW CLOUDBLISS DOES IT

Everything above, codified in five stages: the **Five-M Method**.

STAGE 01	STAGE 02	STAGE 03	STAGE 04	STAGE 05
Map	Measure	Mend	Mechanise	Monitor
Inventory the value streams and follow the real work, not the org chart or the process manual.	Baseline volumes, touch time, cycle time, error and rework rates, so every later claim of benefit is provable.	Eliminate, combine, rearrange and simplify before a single flow is built.	Right-sized automation on Microsoft 365 and the Power Platform, with exceptions, ownership and governance designed in from day one.	Measure against the baseline at 30 and 90 days, then iterate or retire on evidence.

Sometimes stage three is the result. We have ended discovery by recommending a client not automate a process yet, and that report saved them the cost of automating a mistake.

CLOUDBLISS · ON THE MEND STAGE

The common thread is never the industry

We're a UK Microsoft consultancy working with mid-market firms of roughly 50 to 500 people. This method has carried across every domain we've worked in: enterprise asset management operations, finance reporting, PMO delivery. What travels is the discipline of fixing the process before amplifying it.

40⁺

ENGAGEMENTS
DELIVERED

5

STAGES, EVERY
ENGAGEMENT

START HERE

Weighing up what to automate? Start with a short, fixed-scope process discovery.

Or you've automated something and it hasn't paid for itself yet. Either way, the first move is the same: measure before you build.

YOU LEAVE WITH 01

A measured baseline

Volume, touch time, cycle time, error and rework rates, on the processes that matter.

YOU LEAVE WITH 02

A prioritised candidate list

Impact against effort on real numbers, with the tool class matched to each problem.

YOU LEAVE WITH 03

A clear-eyed verdict

What to mend, what to mechanise, and what to leave well alone.

HOW A DISCOVERY RUNS

WEEK 0

Kick-off

Agree the value streams in scope, the people to shadow and the data we can pull.

WEEK 1

Observe and measure

Follow the real work. Sample volumes, touch time, cycle time, errors and rework.

WEEK 2

Report

Baseline, scored candidate list, and a recommendation you can act on or ignore.



TALK TO A FOUNDER, NOT A SALES QUEUE

[Book a discovery call →](#)

Thirty minutes, no pre-sales screening. The quote is the cost.

TRANSFORM · AUTOMATE · GROW
CLOUDBLISS.CO.UK/CONTACT
INFO@CLOUDBLISS.CO.UK