A SHORT BRIEF FOR FOUNDERS AND AGENCY OWNERS

The Payback Gap

Half the budget goes where a demo looks good. The measurable return was somewhere else entirely, and the order you do things in decides whether you can prove any of it.

THE ORDER THAT DECIDES WHETHER IT PAYS BACK

Most failed automation projects did step four first, and skipped the other three entirely.

01
Write it down
So two people can agree what the steps actually are

02
Give it an owner
So somebody is accountable for it working

03
Measure it once
So you have a baseline to compare against

04
Then automate
Now the saving is real, and provable

Skip one to three and there is nothing to compare against.
You cannot prove a saving on a process you never measured before you changed it.

This is why automation budgets get cut in year two.
Not because nothing improved. Because nobody could show that anything had.

Steps one to three are the job I do. Step four is the easy part.

Paul Prado Pacardo
Senior Executive Assistant · Operations and Project Manager

me.kabuhayan.app
linkedin.com/in/paulpacardo
pacardopaul18@gmail.com

WHERE THE RETURN ACTUALLY WAS

The savings came from cancelling invoices, not people.

Budgets concentrate in the visible half of the business. The measurable returns showed up somewhere far less presentable, and they came from a source most owners do not expect.

~50%
of generative AI budget went to sales and marketing
MIT Project NANDA, 2025

Back office
finance, admin, reporting and intake, where the clearer returns appeared
MIT Project NANDA, 2025

External
spend was the source of the saving, not headcount
MIT Project NANDA, 2025

WHY THE SKEW

A pilot in sales photographs well

It demonstrates easily at a board meeting and the story tells itself. Nobody has ever received applause for automating the reconciliation, which is precisely why the reconciliation is still manual.

THE SOURCE

Ending outsourced contracts, cutting agency fees, replacing consultants

MIT found the returns came from reduced external spend rather than reduced headcount. Work the business could suddenly do itself stopped being bought from outside.

WHICH MEANS

The honest version of the argument is also the reassuring one

That is rarer than it sounds. You can tell your team exactly why you are doing this without anyone updating their CV over it.

THE MISMATCH, STATED PLAINLY

The money went where a demo photographs well. The return was in the half nobody wanted to present.

THE TEST, APPLIED

High frequency. Low judgment. Already written down.

Three conditions. A process needs all three. Miss one and the automation either fails or produces a saving nobody can demonstrate.

PAYS BACK

Recurring reporting assembled by hand

Happens weekly, output already defined, nobody has to decide anything. The clearest case there is, and usually the right first build.

PAYS BACK

Intake arriving across four channels

One route in, becoming a tracked item with an owner attached. Removes the checking, not merely the logging.

PAYS BACK

Recurring work creation and scheduling

Predictable, repetitive, and expensive in senior time when somebody rebuilds it by hand every cycle.

NEVER DOES

A judgment call dressed as a workflow

If the answer changes depending on the client, the person or the week, you are automating a decision rather than a process. Decisions do not automate, they just get made worse faster.

NEVER DOES

A process three people disagree about

Automating it does not settle the disagreement. It ships one version faster and louder than the other two, and now the disagreement has a system defending it.

NEVER DOES

Something that happens twice a year

The build, the documentation and the maintenance all cost more than the task itself. Write it down and leave it manual. That is a legitimate answer.

THE WHOLE TEST

High frequency, low judgment, already written down. If a process fails any of the three, fix that before you buy anything.

WHY THE FAILURES FAIL

Nobody automates a process they understand, then regrets it.

The failures share a shape, and it is never the technology. Each one is a decision made before the work was defined.

FAILURE 01

Bought before it was scoped

A tool arrives with a renewal date and no defined first workflow. Six weeks later somebody is hunting for something to point it at, which is the wrong direction entirely.

FAILURE 02

Automated the exception, not the rule

Teams instinctively pick the painful case. Painful usually means unusual, and unusual is exactly what does not repay a build. The boring frequent process is where the hours are.

FAILURE 03

Nobody owned it once it worked

Automations rot. Fields change, people leave, a client asks for something different. Without an owner it fails quietly and everyone reverts to manual without telling anyone.

> "The build is rarely the risk. The definition is.
>
> Which is convenient, because definition is cheap, reversible, and useful even if you never automate a thing.

THE PATTERN UNDERNEATH ALL THREE

Each is an attempt to skip straight to step four. Automation does not create order. It multiplies whatever order already exists, including none.

A NOTE ON THE NUMBERS

Almost every payback figure here is sold by whoever fixes it.

This field has more marketing than evidence. Here is what I checked, what I kept, and what I am refusing to hand you.

What I left out

Figures like "340 percent first-year return" and "4.2 month payback" appear across automation vendors and agencies with no disclosed sample and no stated method. They may be true for somebody. They are not evidence, and quoting them would make this brief an advertisement.

What I kept, with its status attached

MIT Project NANDA for where the return sat, flagged as preliminary and not peer reviewed. Gartner for the agentic forecast, flagged as a prediction rather than a measurement. Consultancy failure rates only with their age and their author's commercial interest stated.

GARTNER

Over 40 percent of agentic AI projects are forecast to be cancelled by the end of 2027

On escalating costs, unclear business value and inadequate risk controls. Published June 2025. This is a forecast, not an observed outcome, and it is frequently re-dated to imply otherwise. Gartner also estimates only around 130 of the thousands of vendors marketing agentic capability genuinely have it.

ERNST AND YOUNG

The widely quoted 30 to 50 percent failure rate for initial RPA projects is about a decade old

Published around 2016 and 2017 by a consultancy that sold RPA implementation at the time. Still cited constantly as though it were current.

WHAT THIS DOES NOT MEAN

None of it means automation does not work. It means the honest evidence is thinner than the marketing, and the difference between success and failure is almost never the tool.

HOW TO DO THE THREE STEPS THAT MATTER

Definition is cheap. It is also most of the value.

Steps one to three take about four weeks, cost nothing but attention, and are worth doing whether or not you ever automate the process.

STEP 01

Write it down from observation

Not a flowchart, and not what people say should happen. A short written account drafted from watching the work actually run. If two people cannot agree on the steps, you have just found the real problem and it was never a tooling problem.

STEP 02

Name one owner

A person, not a team and not a committee. Without this there is nobody to notice when the automation breaks, and it will break the first time an edge case arrives that nobody anticipated.

STEP 03

Measure it once, for four weeks

How long it takes, how often it runs, how often it goes wrong. Change nothing during this period. Skip it and you will never be able to demonstrate that the automation paid for itself.

THEN STEP FOUR

By this point the build is the easy part, and usually the fastest

Most of the value was created in the three steps that felt like they were delaying it. That is the whole argument of this brief.

THE BY-PRODUCT

You now have an SOP whether you automate or not

A written, owned, measured process is a hand-over document, a training aid and a hiring spec. The automation is a bonus on top of something already worth having.

WHY BUDGETS GET CUT IN YEAR TWO

Not because nothing improved. Because nobody measured the before, so nobody could demonstrate the after.

EVIDENCE, NOT ADJECTIVES

I have built the unglamorous half.

Designed, built and maintained by me in a live agency, supporting senior leadership across many concurrent client accounts.

GENERATED

Work that creates itself

A new engagement scaffolds in one click: an epic plus 28 child tickets, phases, components and due dates, identical every time. Recurring sets open two weeks before they are due.

ENFORCED

Standards the system applies

QA checklists attach on ticket creation and reset when work reopens. Incomplete tickets flagged by name to whoever last edited them, in the channel they already use.

DELIVERED

Reporting nobody assembles

Daily stand-ups as direct messages before work starts. Weekly utilization and burn against contracted hours. Intake forms wired straight through to tickets.

SCALE

40+ automation rules, 9 scripted jobs, 8 integrations

Across Jira, Tempo, ScriptRunner and Google Workspace. Two paid automation subscriptions retired and replaced with scripts the team could read and change.

ORDER

Every one of those was documented, owned and measured first

Which is why they are still running, and why the saving could be shown rather than asserted.

CREDENTIALS

Google Project Management Professional Certificate and Google AI Essentials

Plus Anthropic AI Fluency and Model Context Protocol. 31 certifications in total, every one independently verifiable.

> "Paul is dependable, highly organized, and always willing to step in wherever support is needed. His attention to detail, problem-solving skills, and commitment to keeping operations running smoothly made him an invaluable part of the team.
>
> Mary Joy Bertulfo · Executive Assistant · Worked together 2+ years

HOW TO START

Pick the most boring thing you do every week.

Not the one that annoys you most. Annoying usually means unusual, and unusual is exactly what does not pay back. Pick the dull, frequent, unremarkable one.

OPTION 01

Do steps one to three yourself

Write it down, name an owner, measure it for four weeks and change nothing else. Worth doing whether or not you ever automate it, and it costs you attention rather than money.

OPTION 02

Send me the process

I will come back with whether it passes the test, what I would build, and what I would deliberately leave alone. The answer usually includes at least one thing not worth automating.

OPTION 03

Start with the reporting

Recurring reporting assembled by hand is the highest-frequency, lowest-judgment work in most small businesses. It is almost always the right first build.

WHAT I WOULD EXPECT YOU TO FIND

At least one process nobody can describe the same way twice. One that looked automatable and is really a judgment call. And one dull weekly job that has been quietly costing you a day a month.

Available now for remote work

Chief of Staff · Operations · Senior Executive Assistant · Project Manager
Tell me the most boring thing you do weekly and I will tell you whether it is worth automating.

pacardopaul18@gmail.com
linkedin.com/in/paulpacardo
me.kabuhayan.app

Sources: MIT Project NANDA, The GenAI Divide: State of AI in Business, 2025; 300 public deployments reviewed, 52 executive interviews, 153 survey responses; preliminary findings, not peer reviewed. Gartner press release, 25 June 2025, on agentic AI project cancellation; a forecast rather than an observed outcome. Ernst and Young on initial RPA project failure rates, circa 2016 to 2017, published by a consultancy that sold RPA implementation. Deliberately excluded: specific ROI and payback percentages published by automation vendors without disclosed samples or methods. Figures are third-party research, not results I am claiming.