

## INSIGHTS • AI ADOPTION STRATEGY

# Your RPA Graveyard Is Your AI Roadmap

Somewhere on a shared drive, most mid-market companies are sitting on the single most expensive input to AI adoption — and they have it filed under “failed.” The robotic process automation program that stalled out a few years ago left behind a documented inventory of how work actually flows, where it breaks, and where a human has to step in. That inventory is what a multi-week “AI discovery” engagement is built to produce. Before you fund discovery again this budget cycle, it is worth checking whether you already paid for it once.

---

EvoNova Advisors | Insight | July 2026 | 9 min read

## BY THE NUMBERS

**42%**

of companies now abandon most of their AI initiatives before production, up from 17% a year earlier — the same failure that stalled RPA, one technology later.

*S&P Global Market Intelligence, 2025*

**30%**

of failing AI projects will be rescued in 2026 by process intelligence — a documented map of how work actually runs. Exactly what your RPA program produced.

*Forrester, 2025*

**\$1.1B**

spent on process-mining software in 2024, up more than 30% in a year, for tools whose only job is to document how work flows and where it breaks. You already own that map.

*Gartner, 2024*

---

## Everyone is selling discovery. You may have already bought it.

It is budget season, and every mid-market company I talk to is being sold the same first step: an AI discovery engagement. Some number of weeks, some six-figure fee, a team that interviews your people and maps your processes and hands back a prioritized list of where artificial intelligence should go. The pitch is always that you cannot start until you know where to start — and that knowing where to start is the hard, expensive part.

They are right that it is the expensive part. They are wrong that you have never paid for it.

Most companies with more than a few hundred employees ran a robotic process automation program somewhere between 2016 and 2020. A lot of those programs disappointed. The bots were brittle, the maintenance was heavier

than the pitch admitted, and the thing quietly stalled at a dozen automations that break every time a vendor ships a UI update. On the dashboards, it reads as a write-off. The instinct is to treat it as sunk cost and look forward.

That instinct is the mistake. A failed automation program is not a sunk cost. It is a completed discovery phase that you have mislabeled. And this budget cycle, when a CFO is deciding whether to fund an “AI discovery” line item, the honest answer is often that the discovery is already sitting in a folder nobody has opened since the steering committee stopped meeting.

---

## A failed automation program is a completed discovery phase

People forget what RPA actually cost them. The bots were the last thing you built, and the smallest part of the work. Everything expensive happened before a single line of automation ran.

I have written more Process Definition Documents than I care to count, and the pattern never changed. Before anyone automates anything, the standard delivery method makes you document the process at a level of detail no other exercise in the company ever demands. A PDD captures the target process keystroke by keystroke, records the business rules, and enumerates the exceptions — the “happy path” versus every non-standard case. It marks, explicitly, the points where a human has to make a judgment the machine cannot. Good programs then validated those documents with a walkthrough: a subject-matter expert re-performed the process by following the document as written, to surface the edge cases the first draft had missed.

Then you ran the bots, and the logs recorded what happened. How often each process ran, how often it threw an exception, where a person had to step in to keep it moving. Even a program that never scaled produced a real dataset — exception frequency and human-intervention patterns for every process that reached production.

Now set that against what an AI discovery engagement is hired to deliver. A documented inventory of how work flows, where it breaks, and where humans exercise judgment. A ranked view of which processes are painful and high-volume enough to be worth automating. That is the deliverable. That is the PDD library and the opportunity assessment you already commissioned. RPA did more than automate a few tasks; it forced a generation of companies to write down how their operations run, which almost none of them had ever done. When the bots underperformed, the automations got shelved. The documentation stayed true.

*The bots were the last thing you built, and the smallest part of the work.  
Everything expensive happened before a single line of automation ran.*

---

## The bots broke exactly where the value is

The most useful thing about an RPA program is not where it succeeded. It is where it broke.

RPA failed on a predictable set of things, and the failures were structural, not incidental. As one of the field’s early academics put it plainly a decade ago, the technology “can’t structure the data” and “can’t deal with exceptions either” — every time it hit something unstructured or ambiguous, an engineer had to step in and write another rule. Rule-based automation needs the world to hold still: standardized inputs, stable screens, a process that behaves the same way every time. Deloitte’s survey of the era found companies discovering that their processes “are not always well understood, even where robust process documentation exists,” and named process standardization the single biggest obstacle at every stage of maturity. Roughly half of the value in a typical program, McKinsey found, never materialized once teams met the real-world variation the slide deck had ignored.

Read that list of failure points again. It is not a catalog of RPA’s weaknesses so much as a catalog of what large language models are good at. Unstructured documents. Ambiguous inputs. Exceptions that do not fit a rule. Judgment calls that used to require a human. The processes your RPA program rejected as “too complex to automate,” and the ones it automated badly and abandoned, were turned down for the one reason that no longer disqualifies them.

So the exception log is a target list. Every place a bot broke or a human had to intervene is a place where the old technology gave up and the new one begins — and your program already found those places, at your expense, and wrote them down.

EXHIBIT 1

What the graveyard recorded — and why AI wants it

| The artifact RPA left behind | What it documented                                                                                      | Why AI wants it now                                                                              |
|------------------------------|---------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Process Definition Documents | The as-is process at keystroke level, the business rules, and every exception and “happy path” variant. | The workflow map an AI discovery engagement is paid to produce — already written and signed off. |
| Exception & bot-run logs     | How often each process ran and how often it broke — real, measured exception rates, not estimates.      | The target list: precisely where rigid rules failed and language models now win.                 |
| Human-in-the-loop handoffs   | Every step where a person had to exercise judgment the bot could not, marked explicitly.                | The design spec for an AI agent’s guardrails and hand-backs to a human.                          |
| Opportunity assessments      | Process volumes, cycle times, and a pain-ranked backlog of automation candidates.                       | The evidence to sequence the roadmap and clear a CFO’s ROI bar.                                  |

Source: EvoNova Advisors, based on the standard artifacts of an RPA delivery — Process Definition Documents, exception logs, and opportunity assessments.

## You already paid for discovery — once

The reframe matters most when the invoice arrives, so here is the money.

“Discovery” and “assessment” are sold today as packaged products, anywhere from a couple of weeks to a full quarter of work. A scoped AI-readiness or opportunity assessment runs from roughly five to fifty thousand dollars depending on scope; take it to a Big Four firm and it climbs well into six figures. The day-rate math is unforgiving. A major firm’s published government rate card puts a senior partner near ten thousand dollars a day and a junior analyst near three, so a multi-week discovery staffed with a real team clears well into six figures before anyone builds anything. And documenting how processes actually run, from system logs rather than workshop guesses, is now its own billion-dollar software category: process mining generated \$1.1 billion in software revenue in 2024, growing at better than thirty percent a year, with roughly half of large organizations reporting company-wide adoption.

That is the price of discovering how your work flows and where it breaks. Your RPA program produced the same artifact as a byproduct.

What makes this urgent, and not merely clever, is the timing. Three-quarters of CFOs plan to raise technology budgets this year, and nearly half by ten percent or more — but the expectation for headcount growth has collapsed to two percent. Translated: there is money for AI, and no one new to go do the fieldwork. Meanwhile the failure pattern is repeating. More than two in five companies now abandon most of their AI initiatives before production, up sharply in a single year. Gartner’s read on the projects that do succeed is direct — the ones that clear the ROI bar are the ones that integrate AI “into existing workflows and systems,” not the ones that start from a blank page. In that environment, a pre-existing map of your workflows is the difference between the 28 percent that work and the rest.

---

## Yes, it is stale. No, that does not make it worthless.

The argument owes some concessions before the conclusion, because a document written in 2018 is a fair thing to be skeptical of.

First, the documentation has aged, and processes drift. Some of what a PDD recorded is simply wrong now — a system was replaced, a policy changed, a team reorganized. That is real, and it means the documentation is a starting point rather than a finished answer. But refreshing a documented process against current reality is a different exercise from discovering it cold, and a faster and cheaper one. Most of the map still standing beats a blank page, and the least perishable layer — where judgment enters, where exceptions cluster, why the work is hard — tends to survive a system migration intact.

Second, some RPA programs failed because the underlying process was genuinely bad, and no honest reader will let that slide. If a process was a tangle of undocumented variation and manual workarounds, automating it was always going to fail. But that, too, is discovery. Knowing which processes are too broken to automate cleanly is the intelligence that keeps you from pointing an expensive AI agent at them next. The graveyard sorts your processes into the ones worth pursuing and the ones to fix or retire first, and both lists have value.

Third — and this is the concession most consultants will not volunteer — wiring modern AI agents into legacy automation is genuinely hard, and Gartner is blunt that integrating agents into existing systems is technically complex. Fair. But the asset here is not the old bots, and no one should try to resurrect them. It is the *knowledge* the program captured: the process maps, the exception rates, the human-handoff points. You are reusing what you learned, not the code you wrote. The bots can stay buried. The intelligence they generated is what you exhume.

None of these objections closes the case. Grant every one of them in full, and the core fact still stands. The most expensive input to AI adoption is a documented, evidence-based inventory of where your work breaks and where humans intervene — and a stalled RPA program is the only initiative most mid-market companies ever ran that produced exactly that.

---

## Four moves to exhume the roadmap

Turning a graveyard into a roadmap is not a rebuild. It is an afternoon in the archive, followed by a re-triage. Four moves.

**Reclaim the inventory.** Recover the PDDs, opportunity assessments, and bot execution logs from the automation team, the systems integrator, or the vendor who ran the program. If the original center of excellence is long gone, the integrator's deliverables folder and the vendor's implementation records usually still hold copies. This is document retrieval, not a discovery engagement, and it is often the work of a week. What you are assembling is the map you already paid to draw.


**Invert the scorecard.** The processes your RPA program rejected as too complex — too much unstructured data, too many exceptions, too much judgment — are the shortlist now. Take the old “not suitable for automation” pile and re-score it for AI, where ambiguity is a feature rather than a disqualifier.

**Turn the handoff map into the design spec.** Every point where the documentation marked a human intervention is a design decision for an AI agent: where it needs a guardrail, where it hands back to a person, where judgment lives. The human-in-the-loop map your program built defensively, to keep bad automations from doing damage, is the architecture diagram for good ones.

**Sequence by evidence, and fund the one that clears the bar.** You already have the volumes, the cycle times, and the exception rates. Rank the shortlist by hard-dollar impact, pick the process whose documented numbers clear your CFO's ROI test, and fund that — one use case, measured, with a business case built on data you collected rather than assumptions you are paying to generate.

## EXHIBIT 2

## Four moves to turn the archive into a roadmap



Source: EvoNova Advisors.

## Fund the reuse, not the rediscovery

For a leadership team, the uncomfortable part is this: you are about to pay, a second time, to discover something you already documented — and to do it in the budget cycle when every dollar of AI spend faces the hardest ROI scrutiny it has yet seen.

The good news is that the first move is nearly free. Reclaiming your own automation archive costs a few days of someone’s time, commits you to nothing, and tells you within a week whether the discovery a vendor wants to sell you is already sitting on your servers. If it is, you have converted a line item everyone treats as a loss into the cheapest head start available in enterprise AI. The companies that win the next two years will not be the ones that spend the most on discovering where to point AI. They will be the ones who realize they already know — because they wrote it down the last time, and called it a failure. Dust it off.



*An EvoNova Advisors Insight. RPA-era statistics are vintage-flagged and cited deliberately: the argument concerns the legacy automation programs of 2016–2020. AI-adoption, budget, and reuse statistics are current. Directional and vendor-commissioned figures are labeled as such.*

## SOURCES

Deloitte — *The robots are ready. Are you? Untapped advantage in your digital workforce* (Global RPA Survey, n=424) — 2018 (*vintage; RPA-era stall data*)

McKinsey & Company — *How to avoid the three common execution pitfalls that derail automation programs* — 2018 (*vintage*)

McKinsey & Company — *The next acronym you need to know about: RPA* (interview, Leslie Willcocks) — 2016 (*vintage*)

Pega — *Survey: Most Businesses Find RPA Effective But Hard To Deploy and Maintain* (n=509) — 2019 (*directional; vendor-commissioned*)

Forrester Consulting (commissioned by Tricentis) — *Barriers and Best Practices for Scaling RPA* — 2019 (*directional*)

Blue Prism — *Creating a Process Definition Document* (delivery methodology) — 2018 (*standard-practice reference*)

S&P Global Market Intelligence — *Voice of the Enterprise: AI & Machine Learning, Use Cases 2025* (n=1,006) — 2025

Forrester — *Predictions 2026: Automation At The Crossroads* (Leslie Joseph) — 2025

Gartner — *Market Share Analysis: Process Mining Software, Worldwide, 2024* — 2025

Deloitte — *Global Process Mining Survey 2025* (120+ respondents) — 2025

Gartner — *CFOs' Budget Plans Prioritize Growth Functions, Technology and AI in 2026* (300+ CFOs) — 2026

Gartner — *AI Projects in I&O Stall Ahead of Meaningful ROI Returns* (782 I&O leaders) — 2026

Gartner — *Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027* (agentic integration caveat) — 2025

U.S. GSA Advantage — *McKinsey & Company authorized FSS rate card* (Contract 47QRAA22D00E3) — 2024 (*discovery day-rate reference*)

RSM US — *Middle Market AI Survey 2025* (n=966) — 2025

World Economic Forum / National Center for the Middle Market — *AI's mid-market moment* — 2026

---

An EvoNova Advisors Insight, synthesizing published research from Deloitte, McKinsey, Gartner, Forrester, S&P Global, Pega, RSM, the World Economic Forum, and the U.S. GSA. Costs and current-state figures are measured; RPA-era failure statistics are vintage-flagged; directional and vendor-commissioned figures are labeled as such.