

INFOGRAPHIC

How to find & evaluate agentic AI use cases worth building

A practical three-part framework for qualifying automation opportunities — so you build the right things, with the right technology, at the right time.

- 1 Measure the value**
Is there a quantifiable or experiential return worth pursuing?
- 2 Should you automate this?**
Does the process have the volume, frequency, and data readiness?
- 3 Should you use agentic AI?**
Score across five criteria to gauge agent-readiness.

STEP ONE

01 Measure the value

Before anything else, determine whether there is real business value — hard or soft — worth capturing. The most important part of any strategy is what you say no to.

Hard benefits

- Cost takeout — headcount reallocation, reduced manual processing hours
- Cost avoidance — preventing future spend through error reduction
- Revenue impact — faster turnaround times, improved speed to market
- TCO reduction — lower total cost of ownership for core processes
- Risk budget savings — reduction in compliance penalties or losses

Soft benefits

- Operational quality — audit trails, regulatory compliance, accuracy improvements
- Customer experience — CSAT, NPS, faster resolution times
- Employee experience — reduced cognitive load, lower churn in high-stress roles
- Scalability — ability to handle volume growth without proportional headcount
- Life & safety — in clinical or high-risk contexts, agents can surface critical signals

STEP TWO

02 Should you automate this?

Assuming value exists, evaluate whether the process is structurally ready for automation. Don't automate a broken process — fix it first.

CRITERION A

Process volume & frequency

The math is simple: occurrences x handle time = potential time freed. High volume is a clear signal. Low volume with high handle time (e.g. complex loans, investigations) can still qualify.

LOOK FOR

- Strong** High daily/weekly transaction volume across the organization
- Worth it** Low volume but very high handle time (60-90+ min per instance)
- Pause** Quarterly or ad-hoc processes with low total time impact

CRITERION B

Data availability & process readiness

Your automation is only as good as your data foundation. Fragmented data across dozens of unconnected systems means you may need to re-engineer first.

LOOK FOR

- Ready** Data is accessible, consistent, and ideally in one or two systems
- Fixable** Some data gaps, but a clear path to consolidation
- Not yet** Data spread across 10+ silos with no clear ownership or structure

STEP THREE

03 Should you automate this with agentic AI?

Score the process across five weighted criteria. The goal is relative comparison — not a pass/fail threshold. The more processes you score, the clearer your prioritization becomes.


- Logic complexity**
Processes with rigid, well-defined rules (low logic score) are better suited to traditional RPA. Processes with many branches, exceptions, and variable paths are prime candidates for agents. **25%**
- Output variability**
Binary yes/no outputs don't need an agent. Nuanced, contextual, or narrative outputs — summaries, recommendations, custom responses — are where agents excel. **20%**
- Humans in the loop**
High human-touch processes benefit from agents that can interpret free-form human inputs without requiring rigidly structured data — unlike traditional RPA. **20%**
- Flow adaptability**
Unstructured emails, documents in varied formats, or systems that change over time — agents adapt where RPA breaks. This is often the highest-value differentiator for agentic approaches. **25%**
- Integration complexity**
Multiple system integrations add complexity for any automation — but agents can turn that complexity into an advantage, using diverse context to reason and provide better answers. **10%**

Scoring Method

1. Score each criterion 0–100 (multiples of 10)
2. Weighted score = Σ (criterion score x weight)
3. Compare relative scores across your pipeline


There's no hard pass/fail threshold. A score of 21 vs. 76 tells you which process is more agent-ready — use this to prioritize your pipeline, not as a binary gate.

Order entry — parts manufacturer **21.75**
LOW · USE RPA
 Orders via EDI portal with structured inputs; validated against ERP stock and customer records



Verdict: RPA is the right fit. Structured inputs, clear logic, binary outcomes. Agentic AI would be over-engineering.

Foreclosure minimum bid calculation **76.5**
HIGH · USE AGENT
 50+ state/county regulatory rules sourced from public websites; high-risk output reviewed by humans



Verdict: Strong agent candidate. 50+ logic paths, regulations sourced from public web, high-stakes review. Exactly what RPA couldn't handle.

BUILDING YOUR PIPELINE

04 Where are you on the journey?

You can't start at the end. Define your guardrails, build capability, and scale deliberately — this framework only helps you build a rough prioritization.

1 Establish your automation foundation

- Begin with well-defined, high-volume processes using RPA
- Build organizational understanding — not just technical capability
- Set governance guardrails before deploying any AI
- Show the business small wins to build trust and momentum

2 Expand with targeted agentic pilots

- Revisit processes that "hit the cutting room floor" — agents can now tackle them
- Enhance existing automations: inject agents where humans bridge system gaps
- Use the scoring framework to prioritize across your pipeline
- Target processes with unstructured inputs or complex decision logic

1 Reimagine end-to-end with agents in production

- Design processes around agentic capabilities — not legacy constraints
- Use discovery diagnostics to surface opportunities at scale
- Orchestrate agents alongside RPA and human review steps
- Anchor new builds to prior anchors in your scored pipeline