

AI Automation: Process Audit vs Pilot — Which to Start With

A decision guide for choosing your first step in a B2B automation program

Direct answer: should you start with an audit or a pilot?

A process audit **maps** where automation can help — it produces a prioritised list of opportunities and a baseline you can measure against. A pilot **proves** value on one workflow with measurable outcomes. They are complementary, not either/or: a short audit to pick the right first workflow, then a scoped pilot to prove it. Start with an audit if you do not yet know your highest-impact workflow. Start straight into a pilot if you already have an obvious, measurable, rules-based candidate.

What a process audit is — and what it produces

A process audit is a structured look across your teams to find where repetitive, rules-based, or high-volume work is slowing people down. It is diagnostic, not delivery: nothing is automated during an audit. The point is to understand the terrain before you commit effort, so your first build lands on the workflow that actually matters.

A good audit produces:

- **A prioritised list of automation opportunities**, ranked by impact and by how well each fits automation (clear rules, structured inputs, measurable outputs).
- **A baseline for each candidate** — the current volume, handling time, cycle time, and error or rework rate — so you can tell later whether anything actually improved.
- **A shortlist of good first workflows**: the ones that are measurable, rules-based, and contained enough to prove value quickly without touching everything at once.

What a pilot is — and what it produces

A pilot takes one workflow and automates it for real, in a contained scope, measured against the baseline you started with. It is delivery, not analysis: something runs, people use it, and you watch the numbers. The goal is a clear yes/no on whether automation earns its place before you invest in a broader rollout.

A good pilot produces:

- **A working automation on one workflow**, with humans keeping sign-off on any consequential action rather than the system acting unchecked.
- **Measured outcomes against your own baseline** — did handling time, cycle time, or error rate move, and by how much for your team, not a generic benchmark.
- **A decision you can defend**: expand, adjust, or stop — backed by evidence from your own environment, plus an audit trail of what the system did.

Audit vs pilot at a glance

Dimension	Process audit	Pilot
Goal	Map where automation can help and pick the right first workflow.	Prove value on one chosen workflow with measurable outcomes.
Output	A prioritised opportunity list plus a baseline for each candidate.	A working automation on one workflow, measured against that baseline.
Timeframe	Short and diagnostic — typically a matter of weeks, since nothing is built.	Scoped and time-boxed — long enough to gather real usage and results, then review.
Risk	Very low: nothing changes in production; the risk is analysis without action.	Contained: one workflow, human sign-off on consequential actions, an audit trail throughout.
Best when...	You do not yet know your highest-impact workflow, or you have many candidates and no baseline.	You already have an obvious, measurable, rules-based candidate worth proving.

The timeframes above are qualitative on purpose. Real durations depend on the workflow, the data, and your team — treat any universal day-count you see in a vendor pitch with suspicion.

How they fit together: audit → pilot → expand

1. **Audit** to map opportunities and capture a baseline. You come out with a ranked shortlist and the numbers to judge success against.
2. **Pilot** the top candidate. One workflow, human sign-off on consequential actions, an audit trail, and results measured against the baseline from step one.
3. **Expand** only on the strength of that evidence. Move to the next opportunity on the list, repeating the same measured approach rather than automating everything at once.

The core principle

Sequence beats scale. A short audit to choose well, then a scoped pilot to prove value, then expansion earned by evidence — each step de-risks the next.

Avoiding the big-transformation trap

The failure mode to avoid is the sweeping “AI transformation” program — many workflows, many teams, and a long timeline before anyone sees a result. These stall because scope grows faster than value is proven: nothing is measurable, no single owner can point to an outcome, and the effort quietly loses sponsorship.

- **Start narrow, not broad.** One workflow with a measurable baseline beats ten in-flight and none proven.
- **Insist on a baseline.** If you cannot state the before-numbers, you cannot claim an after-improvement — and “it feels faster” is not evidence.
- **Keep humans in control.** Consequential actions go through approval, and every step is logged, so you can trust and audit what ran.
- **Let evidence pull the next step.** Expansion should follow a proven result, not a roadmap drawn before anything shipped.

Which should you choose? A short checklist

Work top to bottom. The first honest “no” tells you to start with an audit; a clean run of “yes” means you are ready to pilot directly.

1. **Do you already know your highest-impact workflow?** If not, audit first — you would be guessing where to build.
2. **Is that workflow rules-based and repetitive?** If the work needs judgement on every case, it is a weaker first pilot.
3. **Can you measure it?** If you cannot state the current volume, handling time, or error rate, an audit to capture a baseline comes first.
4. **Is the scope contained?** One workflow you can pilot without touching everything else is the right size; a sprawling candidate needs narrowing first.
5. **Do you have an owner and a review point?** A pilot needs someone accountable and a date to judge the results against the baseline.

How SG1 approaches this

We favour a scoped pilot measured against your own baseline rather than a universal promise. Consequential actions go through human-in-the-loop approval checkpoints; every AI step is logged in an audit trail; models run on private Azure AI that is never trained on your data; and the system works inside your existing Microsoft 365 tenant rather than copying records out. Where you do not yet know the right first workflow, a short audit to map opportunities and set a baseline comes before that pilot.

We deliberately publish no universal ROI percentage — the only ROI that counts is the one measured against your baseline.