




AN IMPLEMENTATION GUIDE FOR EXECUTIVES

Adopting AI in your business.

A step-by-step framework for a 30-day pilot in a single process — from baseline measurement to the day-30 decision.



What this guide is for

This guide does not explain what AI is. It explains which process to start with in your company, how to set up the pilot, and what to look at on day 30 to make the decision.

Who it is for

The person who makes the decision but will not carry out the work: a general manager, a sales or marketing director, an operations lead. No technical background required, and not a single tool is named.

What is inside

- Process eligibility test — which work to start with, four criteria
- Organizational readiness check — three questions
- Baseline measurement — what to record in the week before the pilot
- The 30-day pilot plan, week by week
- The internal cost of a pilot
- Stopping criteria and the day-30 decision
- A fillable scorecard

How to use it

The scorecard is in the last section. Do not start a pilot before filling it in — the rest of the guide exists so that you can fill it in accurately. Reading takes about 20 minutes; filling in the scorecard may take a few days, because it requires a conversation with the people inside the process.

WHY 30 DAYS

A 90-day plan reads like a project: budget, approval, committee. Thirty days reads like a trial and is easy to start. But the real measure is not the calendar — it is the number of repetitions. No decision should be made before the process has run at least 20 times. For work that runs several times a day, that is a month. For work that runs monthly, it is not.

Why you don't start with the tool

In most companies the AI decision starts like this: someone hears about a tool, a demo is watched, licenses are purchased, the team is informed. Three months later the licenses sit idle and nobody is using them.

This is not because the tool is bad. A tool replaces one step of a process; when no process has been written down, what it replaces stays unclear. The team cannot see where to connect the tool to their own work, so they go back to the method they know. This is not resistance from individuals — it is the natural result of a missing definition.

The order has to be reversed. First you write down how the work is done: which steps, by whom, how often, in how much time. Then you look at which of those steps can be handed to a machine. The tool is chosen last — and at that point the options are usually close to one another. The weight of the decision sits in the process definition, not in the tool.

Not one tool is named in this guide. That is deliberate: for a company moving in the right order, choosing the tool is the easiest step.

The same mistake shows up in automation itself. Automating an inefficient process does not fix it; it simply starts repeating the same error at a speed no human team could manage. AI amplifies this effect, because the output looks correct at first glance.

Which process to start with

Not every process is suited to AI. The candidate process for a first pilot has to meet all four criteria. If one is missing, the pilot stalls on that criterion and you wrongly conclude that AI does not work.

1 • Frequency

The process should run at least 20 times within 30 days. Anything less frequent will not accumulate enough data; you will not be able to see whether anything improved and will have to decide on impressions. Do not run a first pilot on a process that runs less than weekly.

2 • Clear rules

Can the work be described step by step? If the answers are "you learn it with experience" or "it depends," the process is not ready for a pilot. It has to be described first — and that is valuable work regardless of AI.

3 • Accessible data

Does the input sit somewhere: CRM, email, files, a spreadsheet? A process that runs only on knowledge held in people's heads is not suitable for a first pilot.

4 • Low cost of error

The first pilot should sit in a process where a mistake will not harm the business. Work that goes directly to customers, is contractually binding, or produces irreversible output does not come first. As trust is earned, the risk level moves up.

PROCESSES THAT TYPICALLY MEET ALL FOUR

- Preparing the first draft of a proposal
- Classifying an incoming inquiry and routing it to the right person
- Extracting tasks and next steps from meeting notes
- Compiling the weekly report from several sources
- Summarizing long customer correspondence

Do not pick a process that fails one of the four criteria on the assumption that you can push through it. The pilot will jam at exactly that point. What is wrong is the process selection, not AI — and that distinction is easy to miss on day 30.

SECTION 03

Is the company ready

Even when the process is suitable, the company may not be. Three checks:

Is there a process owner

You need one named person, accountable for the outcome, able to give a few hours a week.

"The team will look after it" means nobody owns it. This person should come from the business, not from IT; IT supports the implementation but cannot own the process.

Is the data in good enough shape

If the input is messy (the same customer recorded under three different spellings, half the fields empty), the entire pilot will be spent cleaning data. That is valuable work in itself, but it is not a 30-day pilot, and at the end of it you will have learned nothing about AI.

Can you measure

Can you say how long the process takes today, how many times it runs, and how much rework it produces? If you cannot say that, you will not be able to state the improvement at the end of the pilot either. The next section is exactly for this.

If one of the three is missing, postponing the pilot is the right call. Closing the gap usually takes less time than the pilot itself.

DATA SECURITY AND COMPLIANCE

If the process you pick involves customer data, proposal content, or personal data, run the pilot on an enterprise — privacy-compliant — license. Data entered into free and individual plans can be used to train models, which creates risk under data protection law and under customer contracts alike. License type is the one technical decision that must be settled before the tool is chosen.

Baseline measurement

Before the pilot starts, measure the current process as it stands, for one week. Change nothing. Four numbers are enough.

- **Repetitions** — how many times the process ran during the week
- **Time** — the average minutes per repetition, per person
- **Rework** — how many outputs came back, were redone, or were corrected
- **Who did it** — which role, how many people, how many hours in total

Without these four numbers, all you have at the end of the pilot is an impression. And impressions are almost always favorable: the energy of trying something new pulls the result upward systematically. The baseline is the only protection against that effect.

Who records it

The person inside the process, not the manager. A measurement taken by the manager pushes the person being measured to change their behavior for the duration — and the baseline comes out better than reality. That in turn makes AI look as though it contributed nothing at the end of the pilot.

A PRACTICAL NOTE

Do not build a system to capture the four numbers. Writing them into a single table by hand is enough, and the same table carries on through the pilot. If the measurement tool becomes a project of its own, you will be exhausted before the pilot begins.

30 days

WEEK 1

Setup

The process is written down step by step. The single step to be handed to the machine is chosen — one step, not the whole thing. How the output will be checked is defined: who looks at it, what they look for, how quickly. The tool is chosen this week, not at the start.

WEEK 2

Running in parallel

The new method runs alongside the old one, not instead of it. The two outputs are compared. The goal this week is not speed but reliability: how close is the machine output to the human output, and where does it diverge?

WEEK 3

On its own

Parallel running is removed and the process runs on the new method alone. Human review stays, but nothing is produced twice. Measurement continues without interruption.

WEEK 4

Measure and close out

The same four numbers from the baseline week are collected again. The process owner holds a half-hour conversation with the people doing the work: where did it stick, what did they not trust, what did they find useful.

Three things stay fixed throughout the pilot: the process, the metrics, and the team. Change even one and the number at the end becomes uninterpretable.

SECTION 06

The internal cost of a pilot

This calculation is usually skipped and the pilot is treated as a free trial. The internal load of a typical 30-day pilot:

- **Process owner** — 4–6 hours a week across five weeks including the baseline week: 20–30 hours in total
- **People running the process** — roughly 20–30% extra time each during the parallel week
- **Manager** — a 30-minute review conversation each week: 3 hours in total
- **Tool license** — usually low at pilot scale; the real cost is not here

So the cost of a pilot is attention, not software. If you cannot genuinely set those hours aside, do not start. A half-run pilot leaves behind a false impression that AI does not work, and makes the next attempt inside the company considerably harder.

A REFERENCE POINT

Calculate the current annual cost of the same process: repetitions × average time × the hourly cost of the role involved. The 25–35 hours the pilot takes is usually small next to that figure. If the calculation shows the opposite, the process is not big enough to pilot in the first place.

The human side

Resistance is usually aimed at uncertainty rather than technology. If two questions go unanswered, the team will not reject the pilot openly — it will quietly stop working.

"Will this take my job?"

Answer honestly. If one step of the work is going to the machine, say where the freed-up time will go. If you cannot say, the team assumes the worst, and no pilot survives that assumption. Put the answer in writing; a verbal reassurance does not leave the meeting room.

"Who is accountable if the output is wrong?"

During the pilot, accountability for the output sits with a person: the machine drafts, a person approves, and responsibility stays with the approver. If this rule is not written down, nobody dares approve anything and the process drifts back to its old form on its own.

Who to start with

Start the pilot with a volunteer — the most curious person, not the most skeptical one. The skeptic is far more valuable on day 30, when the result is being assessed; that is the moment they ask the questions nobody else is asking.

SECTION 08

The day-30 decision

Write the decision criteria at the **start** of the pilot, not at the end. Criteria written at the end always fit whatever happened.

SCALE	There is meaningful improvement in the measured time (set the threshold yourself up front; 20% is a reasonable starting point), the rework rate is no worse than the baseline, and the people running the process want to continue. All three have to hold.
FIX	There is improvement, but the rework rate rose, or the team does not trust the output. The problem is almost always in the review step, not the tool. Redefine the review and repeat the 30 days.
STOP	There is no improvement against the baseline, or the process never settled. Stopping is not failure: learning in 30 days that a process is unsuitable is better than paying for an unused license for two years.

Make stopping legitimate from the outset. A pilot that cannot be stopped is not a pilot — it is an investment with no end date.

SECTION 09

Scorecard

0 = no · 1 = partly · 2 = yes. Fill it in by talking to the people inside the process, not by guessing.

CRITERION	0	1	2
A · PROCESS ELIGIBILITY			
Frequency The process runs at least 20 times in 30 days	☐	☐	☐
Clear rules How the work is done can be described step by step	☐	☐	☐
Accessible data The input sits in a system, not only in people's heads	☐	☐	☐
Low cost of error A faulty output is recoverable and does not go straight to the customer	☐	☐	☐
B · ORGANIZATIONAL READINESS			
Process owner One named person, 4–6 hours a week, from the business	☐	☐	☐
Data quality The input will not require a cleanup project first	☐	☐	☐
Measurement Time, repetitions and rework can be recorded today	☐	☐	☐
Total score	/ 14		

11–14 **You are ready to pilot.** Start with Section 4, then follow Section 5.

7–10 **One foundation is missing.** Close your lowest-scoring row first.

0–6 **The problem is not AI.** Process definition and measurement come first.

Scored below 10? The order in which to close the gaps is usually clearer in conversation than in a guide. A 30-minute diagnostic call: eidosgrowth.com/en/contact

The five most common mistakes

- 1 Starting on several processes at once.** None of them settles, you cannot tell which one worked, and because attention is split three ways all three stay half-finished.
- 2 Skipping the baseline measurement.** On day 30 you have nothing to say beyond "it felt better" — and you do not take that sentence to a board.
- 3 Choosing the tool first.** The chosen tool shapes the process around its own logic, when the process should be shaping the tool.
- 4 Handing the pilot to IT.** The owner is the person who knows the work. IT supports the implementation but cannot own the process, because it cannot judge whether the output is correct.
- 5 Not writing the stopping criterion.** A pilot without one stays "in progress" forever and nobody asks for the result.

LAST WORD

What comes next

This guide was written to help you set up a single pilot properly. Scaling — moving to a second and third process, connecting processes to one another, making the measurement permanent, reshaping how the team works around it — is a different job, and it was deliberately left outside this guide.

But at the end of the 30 days you will have one thing: a decision about a single process grounded in measurement rather than impression. In most companies that is precisely what is missing.

Let's talk about your score

If you have filled in the scorecard, we can clarify in 30 minutes which gaps to close in which order and which process makes the most sense to start with. This is not a pitch; we look at where your system is stuck.

Book a call	eidosgrowth.com/en/contact
Email	hello@eidosgrowth.com
LinkedIn	linkedin.com/company/eidosgrowth

Eidos Growth builds growth, positioning and sales systems for companies selling complex products and services. It treats AI not as a separate project but as part of those systems.

© 2026 Eidos Growth · eidosgrowth.com