Quether

GUIDE

Bringing AI into a mid-sized company

The order of decisions, from a business pain point to something that stays up and running

FORMAT

PDF guide, with the prioritisation matrix and the preliminary checklist

WHO IT'S FOR

Managing directors, operations directors or whoever has to answer for this being of any use, in a company that has no data department.

WHAT IT TAKES

A twenty-minute read. The pilot it proposes runs for ninety days.

LAST REVIEWED

26 August 2026

https://quether.com/en/resources/implementing-ai-in-a-mid-sized-company

What you take away

A method guide for companies turning over €3–50M that have already tried a tool and haven't managed to get anything into production. Where to start, how to prioritise, what a pilot that works looks like and what has to happen for it to survive.

+ Why pilots die, and the only order of decisions that prevents it
+ The matrix for choosing the first use case: value, effort, risk and who sustains it
+ The four questions about the data that you need to answer before looking at a tool
+ The ninety-day pilot, week by week, with its stop criterion
+ The five conditions for moving from pilot to production without it collapsing within six months

Why pilots die

The sequence repeats itself in almost every mid-sized company. Someone tries a tool and it works for them. It gets bought for the team. For three weeks there's enthusiasm. By the third month two people are using it, nobody can say what has improved and the conversation has moved on to the next tool.

It isn't a technology problem: the tool did what it promised. It's a problem of order. It started with the answer—the tool—without having written down the question—what hurts, and how much. And an answer without a question can't be evaluated, so it can't be defended when the first busy week comes along.

> **One technology changes a task. Several together change the system.**
>
> That's why the first use case isn't chosen for how impressive it is, but for how well it can be measured.

The right order, in four steps

It isn't negotiable and it can't be done in reverse. Each step conditions the next, and skipping one explains almost every project that never gets there.

From nuisance to system

Four decisions, in this order. If you catch yourself talking about a tool before step 4, you're back at the start.

01

The pain point, with its number

CLOSES WITH

"It costs us _____, and today that's ___ hours a month or €___ a year." If you can't put a figure on it, it isn't a case yet: it's a nuisance.

02

The use case, narrowed down

CLOSES WITH

One specific task, for one specific person, in one specific process. "Improve customer service" isn't a use case; "draft the first reply to a warranty claim" is.

03

The data it needs

CLOSES WITH

Where it lives, who maintains it, what state it's in and who can authorise its use. Without those four answers, you don't go on.

04

The tool

CLOSES WITH

It's chosen last, and it's chosen against the three previous steps. It's the easiest decision to change, and that's why it comes last: changing it doesn't invalidate the work done.

The test of whether step 1 is done properly

Show the sentence to the person who does that work today. If they say "yes, exactly that", the case exists. If they say "well, it depends", it still needs narrowing down. Nobody knows the real process better than the person who puts up with it every day, and they're almost never in the meeting where it's decided.

Where the first case usually is

In a mid-sized company, the value is rarely where people talk about it most. These four places share one feature: the work is already written down somewhere, so it can be measured before and after.

STRUCTURE

The knowledge only one person has

Quotes, cost breakdowns, acceptance criteria, technical answers. When someone has been there twenty years and everything goes through them, the risk isn't one of efficiency: it's one of continuity. Capturing that judgement in something people can consult is worth it on its own, with or without AI.

SALES AND AFTER-SALES

The repeated answer

Offers, tender documents, complaints, the usual questions. The first draft is where the time goes; the review is still human and should stay that way. This one is easy to measure: time to first draft and number of rounds.

OPERATIONS

The data typed in twice

From the delivery note to the ERP, from the email to the CRM, from the job sheet to the timesheet. It isn't the flashiest case and it's usually the most profitable, because the error it removes costs more than the minute it saves.

ADMINISTRATION

Reading other people's documents

Contracts, supplier invoices, regulations, tender specifications. Slow work, subject to deadlines and with consequences if something slips through. Start by extracting and comparing, never by deciding.

What the first case shouldn't be

Nothing that makes decisions about people—recruitment, appraisal, promotion, shifts—and nothing that talks directly to the customer without review. The first enters the territory with the highest regulatory demands; the second puts your brand in the hands of a system you don't know yet. Both will come; neither is where you start.

How to choose between several candidates

Score each candidate from 1 to 5 in the five columns and rank them by the total. The fifth column is the one that has saved the most pilots and the one almost nobody includes.

CRITERION	WHAT YOU ASK	SCORES HIGH IF…
Value	How much is solving it worth, in hours or in euros a year?	The figure from step 1 is large and someone in finance recognises it.
Measurability	Can you measure the before and after without setting up anything new?	The data is already recorded today, even if by hand.
Effort	What does it take to have something working in ninety days?	It doesn't require integrating systems or migrating data to get started.
Risk	What happens if the result is bad and nobody notices?	The error is visible, reversible and a person reviews it before it goes out.
Owner	Who answers for this still being used six months from now?	They have a name, they're inside the process and it's in their interest for it to work.

The owner column isn't a formality. A case without an internal owner holds up while the project's enthusiasm lasts and falls apart exactly when the project ends. If nobody inside wants to be the owner, that's information about the case, not about the people.

Data is where projects stall

Almost no implementation stalls because of the model. They stall because the data needed was in five places, under three different names, maintained by nobody in particular and with a part that couldn't be used. These four questions are answered in writing before looking at any tool.

01 Where does it live? A system, a server, a shared folder, someone's head. Every answer is valid except “I don't know”.

02 Who maintains it? If the answer is “everyone”, nobody maintains it and its quality is a lottery.

03 What state is it in? How much is duplicated, how much is out of date and which fields are empty in exactly the rows that matter.

04 Can it be used? Personal data, customer information covered by a contract, supplier material under confidentiality. The answer comes from whoever can authorise it, not from whoever wants to use it.

Where the information you give a tool ends up

Before uploading a single document: where it's stored, for how long, whether it's used to train models, who at the provider can see it and what happens when you cancel the contract. It's a five-minute conversation with the provider and their answer should be in writing. If they avoid it, you've already learnt something.

The ninety-day pilot

Ninety days are enough to know whether something works and too few for it to turn into an endless project. The end date is set on day one, and with it the stop criterion.

Ninety days

One case, one person, one number. Not two cases “while we're at it”: the second case doubles the work and splits attention just when it's needed in full.

W1-2

The baseline

CLOSES WITH

Today's number, actually measured and not estimated in a meeting: how long it takes, how many rounds it goes through, how many errors slip in. Without a baseline there's no improvement, just opinions.

W3-6

The minimum version, in one person's hands

CLOSES WITH

Something the person who does that work actually uses in their day. With mandatory human review and an obvious place to say “this came out wrong”.

W7-10

Adjusting with real use

CLOSES WITH

The corrections that come from use, not from a list of improvements written in advance. This is where you discover what the process was doing that nobody had mentioned.

W11-12 **The decision**

CLOSES WITH

The same number as in week 1, measured the same way. And one of three decisions, in writing: it's rolled out further, it's adjusted and extended once, or it's stopped. Stopping well is a result.

The stop criterion is written on day one

What result would make us drop it, decided before there are any results. Written afterwards, there's always a reason to carry on for another quarter: it's the most expensive bias in any pilot project and the only one that can be prevented with a sentence.

From pilot to production: five conditions

A pilot that works isn't a system that stays. These five conditions separate one from the other, and all five are about organisation, not technology.

01 It sits inside the process, not alongside it. If people have to step out of their usual workflow to use it, they'll stop using it as soon as there's a rush.

02 It has a named owner, and that name is in their job description, not in the minutes of a meeting.

03 It has human review at the point where an error would be costly, and that review leaves a trail of who reviewed what.

04 It's measured with the same number as the baseline, on a dashboard someone looks at every month without being reminded.

05 It's documented well enough for someone else to pick it up: what it does, with what data, what to do when it fails and who to tell.

The fifth condition is the one most often skipped and the one that turns an improvement into a dependency. A system that only the person who built it understands isn't a company asset: it's a risk with a friendly face.

People: what decides whether this stays

The resistance is almost never to the tool. It's to the implicit message, which people read before anyone says it: if this does my job, what am I here for? As long as that

question has no public answer, everything else is theatre. And the answer has to be true, because it will be checked.

Training isn't a generic course for the whole workforce. It's matching what each person needs to know to what they actually do: who uses the system, who reviews its output, who makes decisions with it and who answers if it goes wrong. As well as being what works, it's what European law requires, and since earlier than almost everyone thinks.

LEVEL 1

The person who uses it

What it's for, what it isn't for, how to recognise bad output and what to do with it. With examples from their own work, not generic ones.

LEVEL 2

The person who reviews it

What to look at and in what order, where the system tends to fail and the threshold beyond which things get escalated. Review without judgement turns into a click.

LEVEL 3

The person who decides

What can and can't be claimed with this, which obligations apply and what the company is liable for. It's the shortest training and the one most often skipped.

The regulatory minimum, said early

Two obligations under the EU AI Act affect any company using these tools today, whatever its size: AI literacy for the workforce, in force since February 2025, and transparency, which requires telling people when they're interacting with an AI and labelling certain synthetic content.

It isn't a layer added at the end. Deciding where human review comes in, what data can be used and what the customer is told are design decisions for the use case: taking them at step 2 costs a conversation; taking them once it's in production costs redoing it.

The detail is in the other guide

What exactly it requires, of whom, since when and what was postponed by the digital omnibus of July 2026 is in “The EU AI Act, explained for leadership”. It lives in one place on purpose: it's what changes most and it can't be duplicated.

Twelve questions before signing anything

- [ ] What specific pain point does it solve, and how much is that pain worth a year?
- [ ] What exact task, for which person, in which process?
- [ ] What's today's number and who actually measured it?
- [ ] Where does the data it needs live and who maintains it?
- [ ] Who can authorise its use, and have they authorised it in writing?
- [ ] Where is what we upload stored, for how long, and is it used for training?
- [ ] What happens to our data when we cancel the contract?
- [ ] At what point does a person step in to review, and is there a record of it?
- [ ] Who is the internal owner, by name?
- [ ] What result would make us stop it, written down before we start?
- [ ] Who needs to be trained, and in what exactly, according to their role?
- [ ] Does any transparency or high-risk obligation apply to this case?

Quether

Quether works alongside leaders of mid-sized companies to transform their organisation—and themselves—without losing their essence.

This material, always up to date, at https://quether.com/en/resources/implementing-ai-in-a-mid-sized-company