

AI TRANSFORMATION PLAYBOOK

The AI Transformation Playbook

How to Pilot, Prove, and Scale AI Across Your Organization

.TrackImpact



CONTENTS

In This Guide

1 — The Multiplier Problem	3
2 — Build the Case	5
3 — Run the Pilot	8
4 — Track & Evidence	12
5 — Report & Scale	14
6 — Put It Into Action	15

You Don't Move Your Own Numbers. You Move Everyone Else's.

Most guidance on using AI well is written for one person trying to get better at their own job. That's not your job. You're the one who has to make AI adoption real across a team, a department, or an entire company — and you're judged on outcomes you don't personally produce.



That's the multiplier problem. An individual contributor who gets great at AI improves one person's output. A transformation lead, an AI officer, a "make this real" owner — whatever the title — is supposed to improve everyone's. The tools for doing that well are different: less about prompting technique, more about picking the right starting point, getting a team to actually adopt something, and being able to prove, in numbers a skeptical CFO will accept, that it worked.

You're also usually the only person in the building whose job depends on this succeeding company-wide. Individual teams have day jobs; AI adoption is a project layered on top of them. Without someone owning the rollout, the pattern is completely predictable: a few enthusiastic early adopters use AI well, most people don't touch it, and six months later there's no organization-wide answer to "what did we actually get for this."

This guide is the playbook for avoiding that outcome — building the case, running a real pilot, generating evidence as you go instead of scrambling for it later, and using that evidence to earn the budget and mandate to scale. It assumes you're incentivized to get this in front of as many people in your organization as possible, because that's genuinely how the math works in your favor: every team you bring in is more adoption, more evidence, and a stronger case for the next one.

Where AI Actually Delivers

✓ **Org-Wide Visibility**

✓ **A Repeatable Rollout**

✓ **Evidence for Budget**

Where It Still Needs a Human

✗ **Team-by-Team Judgment Calls**

✗ **Political Buy-In**

What Leadership Actually Asks Before Funding a Pilot

Before you get a budget line or a team's time, someone with authority is going to ask a short, predictable set of questions. Answering them clearly, in advance, is most of what separates a pilot that gets approved from one that dies in a hallway conversation.



The Questions You'll Actually Get

Not "is AI important" — everyone already agrees on that. The real questions are narrower: which team, doing what, for how long, measured how, and who's accountable if it doesn't work. A pitch that answers those five things in a page beats a slide deck that argues AI is the future.

Picking Where to Start

Resist the instinct to pick the hardest, highest-stakes function first to prove a point. The strongest first pilot has three things: a function with real, repetitive task-level work (not just "strategy"), a manager who's willing to champion it rather than tolerate it, and a task that's easy to measure — time spent, volume handled, error rate — so the eventual report isn't an argument about what counts as success.

The Pitch, in One Page

Problem: Name the specific, current cost — hours spent on a task, inconsistency across the team, work that doesn't scale with headcount. Not "we're behind on AI."

Pilot Scope: One team or function, a defined set of tasks, a fixed window — 4 to 8 weeks is usually enough to generate a real read without dragging on.

Success Metric: Decided before the pilot starts, not after — time saved, adoption rate, or output volume, tied to something the team already tracks.

Timeline: Setup week, pilot window, report-back date. Put the report-back date on the calendar now, not after you see the results.

The Ask: What you actually need — usually just the pilot team's time and a green light, not new budget. Save the budget ask for the report-back, once you have numbers.


The honest version of this section: a pilot with a vague success metric doesn't fail, it just becomes unfalsifiable — nobody can tell if it worked, so nothing changes and nobody gets blamed. Deciding the metric before you start is what turns "it went fine" into a real yes-or-no answer.

BEFORE YOU START

What TrackImpact Actually Is

TrackImpact is a live platform for exactly this: job profiles matched to real tasks, tool-stack scoping, individual usage logging, and org-wide reporting, all in one dashboard. It's what turns the pilot in this guide from a spreadsheet exercise into something your team logs as they go and leadership can see in real time.

The panel below reflects the real admin dashboard — the same example figures shown on trackimpact.io, and the same kind of view you and your team will use to run the pilot in the rest of this guide.



See TrackImpact → www.trackimpact.io/

Two Moves Before Anyone Touches a Tool

The pilot itself runs on a five-step loop — get the profiles, set the stack, log the usage, roll it up, report the impact. This section covers the first two, which are entirely on you before the team does anything differently.



The Loop, End to End

#1

Get the Profiles

Every role on the pilot team gets a TrackImpact profile showing exactly where AI is already effective for that job.

#2

Set the Stack

You tell TrackImpact which AI tools are actually approved and licensed. Every recommendation is scoped to that list from then on.

#3

Log the Usage

As the team works, they log real activity through the TrackImpact buttons in their profile and dashboard — seconds per task, not a survey.

#4

Roll It Up

The dashboard aggregates logged activity by team and department automatically — nobody manually compiles anything.

#5

Report the Impact

TrackImpact turns the aggregated activity into a report you can actually bring to leadership.

TRACKIMPACT — SEARCH PROFILES

Not Sure Where to Start? Search Your Job Title.

TrackImpact's marketplace covers thousands of O*NET job titles, and every profile breaks the role down into its real, recurring tasks — not "use AI for marketing," but the actual activities someone in that seat does every week — each matched to a concrete AI approach and an automation-readiness read.

Instead of guessing which tasks are worth automating, search the job title and see the answers already mapped out.

SEARCH

"Senior Financial Analyst"

TASKS MAPPED

38 recurring tasks found

→ SAMPLE TASK MATCH

"Reconcile the variance report against last month, flag anything outside a 5% band."

HITL SCORE: MEDIUM — AI DRAFTS, YOU REVIEW

Search your job → www.trackimpact.io/search/

TRACKIMPACT — SET THE STACK

You Decide What's Approved. Every Profile Respects It.

Before anyone on the team logs a single task, you tell TrackImpact which AI tools are actually licensed and approved — function by function. From that point on, every recommendation any profile surfaces is scoped to that list. Nobody gets pointed at a tool they can't open, and nothing your company has already paid for goes unused because no one knew to suggest it.

It's a five-minute setup that quietly prevents the single most common way pilots stall: a team member trying to follow guidance built around a tool they don't have access to.

YOUR APPROVED STACK

Sales & CS HubSpot AI, Otter.ai

Finance & Ops Excel Copilot

Marketing Jasper, ChatGPT
Enterprise

→ EVERY RECOMMENDATION IS SCOPED TO IT

*“Reconcile the variance report against last month” —
matched only to tools your team can actually open, never
a generic suggestion for something you haven't licensed.*

ONE APPROVED LIST, EVERY PROFILE

See how it works → www.trackimpact.io/

Turning Real Usage Into a Record Leadership Trusts

This is the step that turns a pilot from a vibe into a number. Everything downstream — the report, the budget conversation, the case for the next team — depends on the logging actually happening, and happening honestly.



What Logging Looks Like Today

Once tasks and tools are set up, each person on the pilot logs their real usage through the TrackImpact buttons built into their profile and their dashboard — a few seconds per task: what they did, which tool, roughly how long it took or saved. That's deliberately low-friction. A logging process that takes longer than the task itself doesn't survive past week one.

To be precise about what's here today: job-profile matching, tool-stack scoping, and usage logging via the TrackImpact buttons in each profile and dashboard are real, live features your pilot uses right now — not a roadmap pitch. The record it builds is accurate and honestly-sourced because your team logs it themselves, a few seconds at a time.

Keeping It Honest


Two things make the difference between a record leadership trusts and one they quietly discount: consistency (the same task logged the same way by everyone, not just the enthusiastic adopters) and including the misses, not just the wins — a task where AI didn't help is still useful data, and a report with zero negative results reads as filtered, not successful.

TRACKIMPACT — TEAMS

See Impact Across the Whole Organization

Teams brings every individual's tracked and scored activity together at the organizational level — so leadership can see adoption and value by department, site, or role, and compare who's driving results.

It's the difference between a handful of enthusiastic early adopters and a rollout that can actually be managed and reported on.



See how it works → trackimpact.io

Turning the Record Into the Next Budget Conversation

The report-back date you put on the calendar in Section 2 has arrived. This is where logged, aggregated activity becomes the case for expanding past the pilot team.



What a Report That Lands Actually Has

The metric you committed to in Section 2, stated plainly — not reframed after the fact to flatter whatever the number turned out to be. The adoption rate across the pilot team, not just the top performer. At least one concrete, specific example of a task that changed, described in enough detail that a skeptical listener believes it really happened. And an honest account of what didn't work, which is what makes the rest of the report credible rather than a sales pitch.

Making the Scale Case

The strongest ask isn't "give us more budget for AI" — it's "here's what this specific, measured pilot returned, here's the team we'd add next and why, and here's what the same report will show for them." A transformation lead who can point to one real, evidenced rollout has a fundamentally stronger position than one still arguing in hypotheticals, because the next conversation isn't about whether this works — it's about which team gets it next.

A Note on Pace

Scaling team by team, with a real report between each one, is slower than announcing a company-wide rollout on day one — and it's also the version that's still running in a year. Each report is both evidence for the next expansion and a check that the process still works at the next size up, before you've committed the whole organization to it.

Five Questions Before You Start

Answer these before you write the pitch in Section 2 — they're the inputs everything else in this guide assumes you already have.

Which function has real, repetitive, task-level work and a manager willing to champion this?

e.g. customer support — first-response drafting

What's the one metric that decides whether the pilot worked?

e.g. average first-response time, tracked weekly

Which AI tools are already licensed and approved for that function?

e.g. Zendesk AI, ChatGPT Enterprise

What's the report-back date, and who's it going to?

e.g. 6 weeks out, to the VP of Customer Success

If it works, which function gets the pitch next — and why that one?

e.g. finance — similar task-level repetition, different exec sponsor

Get Started

The fastest way to start is exactly where Section 3 begins: get profiles in front of the pilot team, then set your approved stack.

Find your team's roles: trackimpact.io/search

Start tracking: [your TrackImpact dashboard](#) — [Track Impact](#)



Ade Molajo

Founder, TrackImpact

LinkedIn: linkedin.com/in/ademolajo

Email: ade.molajo@trackimpact.io

Web: trackimpact.io

TRACKIMPACT — TEAM ONBOARDING

Want This Running Across Your Org in Weeks, Not Quarters?

Everything in this guide is doable on your own. If you'd rather not run the rollout solo, TrackImpact's team will map every function, set the approved stack department by department, train the people who'll actually log usage, and get your first reporting cycle in front of leadership — done with your team, not handed to it.

It's built for transformation leads who already have executive air cover and need the pilot live fast, not a slow-burn DIY project competing with everything else on your plate.

WEEK 1 Map every team and department into TrackImpact.

WEEK 2 Set the approved AI stack, function by function.

WEEKS 3-4 First reporting cycle live, ready for leadership.

Watch how it works → www.trackimpact.io/lp/teams/

Photo Credits

All photography sourced from Unsplash. The Unsplash License does not require attribution, but it's listed here as good practice.

- Small team in a candid discussion around a desk — Austin Distel / Unsplash (img-transform-framing.jpg)
- Team reviewing plans at a whiteboard, casual office — Austin Distel / Unsplash (img-transform-case.jpg)
- Diverse team collaborating around a laptop in a glass-walled office — Vitaly Gariev / Unsplash (img-transform-pilot.jpg)
- Person focused on laptop work at a desk — Vitaly Gariev / Unsplash (img-transform-track.jpg)
- Team reviewing plans together in a sunlit office — Sable Flow / Unsplash (img-transform-scale.jpg)
- Team lead presenting to colleagues in a casual office meeting — Austin Distel / Unsplash (img-transform-cover.jpg)