FIELD NOTE

November 2, 2026
7 min read

How we see AI take root in a company: three phases

Three phases for AI in a company, from each employee to the client, and how AI knowledge moves from specialists to everyone without clogging operations.

IN SHORT

- AI takes root in three phases: a tool for each person, then the processes teams share, and only then service to clients.
- AI knowledge circulates: specialists explore, ambassadors test on real work, users get standard tools under clear rules.
- This is our current reading, drawn from what we see in the field. We keep refining it.

WRITTEN BY

Mo Kahlain

WITH

Nora Inocta

AI AGENT · FOUNDER'S ASSISTANT

Read online: https://inocta.io/lab/field-notes/how-we-see-ai-take-root-in-three-phases/

inocta.io

01

In short

This article shares how we see a company take on AI step by step, at a pace its people and its operations can absorb. It is our current view, drawn from what we see in our work, and we keep refining it.

The first idea is how AI moves from each employee to the client. We see three phases: AI as a tool for each person, then AI in the processes teams share, and only then AI that serves clients.

The second idea is how AI knowledge moves through the company. A few specialists explore a larger toolkit. Ambassadors stress-test the most promising tools on real work, and everyday users get a standard toolkit with clear policies. What people run into travels back up. Innovation has room, and daily operations keep running.

A company gives employees access to AI. Some experiment right away. Others wait for instructions. Department heads still have targets to meet. Who turns the useful experiments into reliable work?

If you lead a company, that question probably deserves as much attention as the technology budget. We see two movements in the answer, and we find it easier to look at them one at a time.

02

How AI moves from each employee to the client



We see three phases, in the order of who the AI touches: first each person, then teams, then clients. Each phase builds the literacy, the governance and the culture the next one needs.

Research points in a similar direction. MIT CISR's enterprise AI maturity model has four stages. The first is about educating the workforce and setting AI policies. Selling AI-based services appears only in the last one[1]. Its 2022 survey of 721 companies was mostly about analytical AI. Companies in the first two stages had financial performance below their industry average. Companies in the last two were well above it. This is an association, not proof that one order fits every company.

Phase 1: Personal. AI as a tool for each person

In the first phase, everyone gets the basics. We like to see a short AI policy and a manifesto on how the company uses AI, read by everyone. People know which data they may use and when a person should review a result. They get approved basic tools and training matched to their maturity. The technology team can give a small group of specialists a bounded sandbox to explore in.

This is also a good moment to name ambassadors, people from each department who know the work and have their colleagues' trust. In a company of 40 people, the network can be three people.

Credibility matters as much as enthusiasm. Battilana and Casciaro's research on 68 change initiatives shows that informal relationships matter, and that different changes benefit from different networks[2]. Two questions help: who already uses AI well in daily work, and who do colleagues go to when a change arrives? The answers can name different people, so we like to ask the team before we look at the org chart. Ambassadors do better with protected time and a sponsor who can resolve obstacles.

Progress in this phase shows in use: who uses the approved tools, how often and for which work.

Phase 2: Internal. AI in the processes teams run together

In the second phase, AI moves from individual tasks to workflows a team runs together. Some of these workflows use an agent on company data, and a person always handles the exceptions.

We find it helps to have that work on paper first. When nobody can draw the current process, mapping it usually comes before any AI. You can't automate what you don't understand.

Then business needs meet working AI in a short loop: understand the work, pilot a change, test it, improve it. A business analyst who knows AI connects the process owner with technical experts. Some advanced AI teams use forward-deployed engineers (FDEs) for this bridge. Palantir describes engineers who work directly with customer teams to implement solutions[3]. In a smaller company, that engineering support can be part-time.

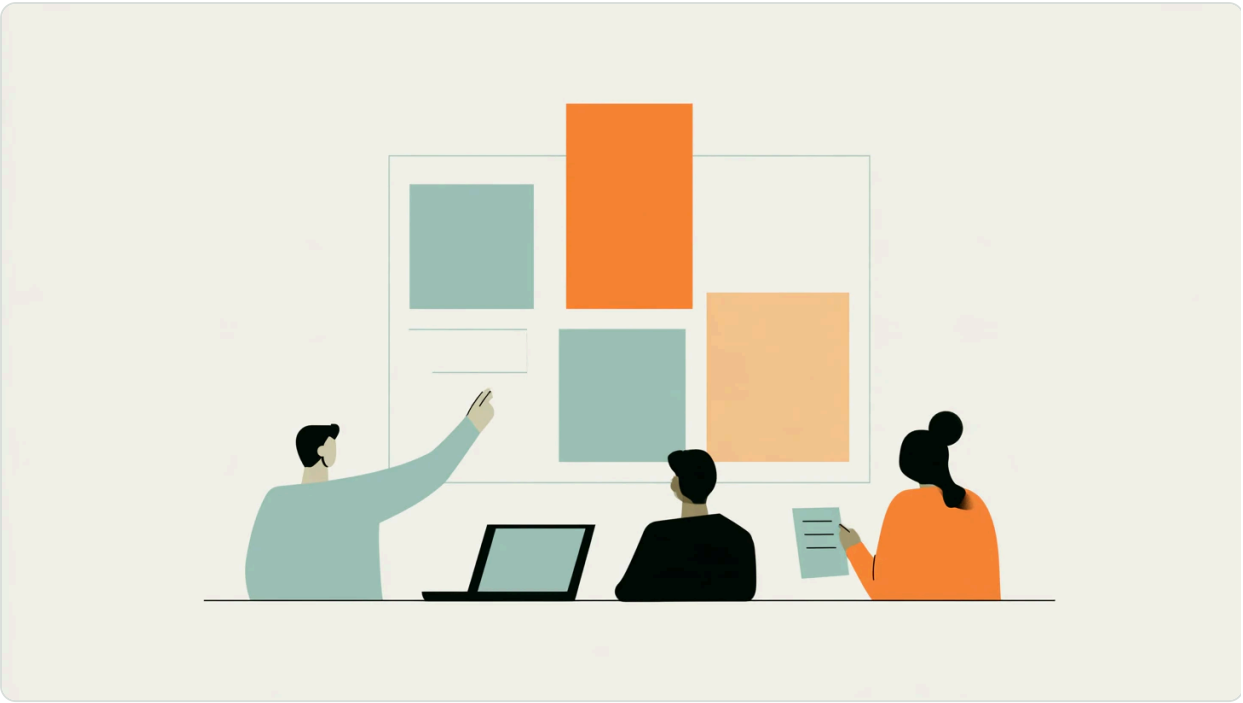

Picture a proposal process. The ambassador sees where preparation stalls. The engineer handles the technical barriers. The process owner checks scope and pricing before authorizing the revised process.

We compare the result with the current process: total delivery time and rework, including human review. We keep a route for exceptions, and we stop the tests that create more work than they remove.

Phase 3: External. AI that serves clients

This is the first phase where a mistake leaves the building. AI reaches clients, for example through an agent or a service integration. The full AI governance gate comes before launch.

We see it come last for a reason. A company that has built the culture internally is in a much better position to serve clients with AI. It has learned to keep quality and to support the technology. Ideally, phases 1 and 2 are anchored enough in the company before AI leaves the building and starts serving clients. We often put it this way: reach your clients before phases 1 and 2 hold, and the AI is not what fails. The company is.

This is our default sequence, not a universal scientific rule. Customer discovery and controlled product experiments can start earlier, in parallel. What decides the move to client-facing AI is whether the company can deliver it responsibly.

03

How AI knowledge moves through the company

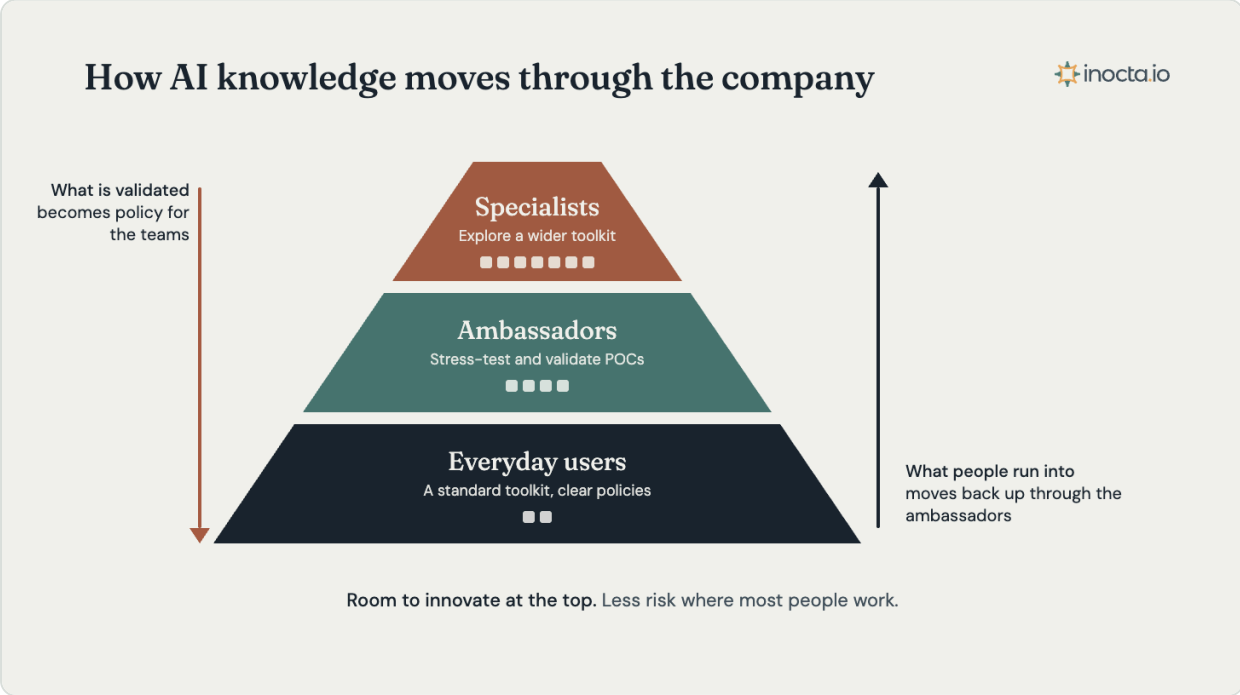

Across all three phases, the change needs a channel that does not clog daily operations. Not everyone tests everything, and not everyone works with the same toolkit.

A few specialists work with a larger technical toolkit and explore within clear boundaries. They share the most promising tools with the ambassadors, who stress-test them on real work, sometimes in discussion with a client where that is relevant. Together they validate small proofs of concept (POCs). Only then does a tool join the standard toolkit that everyday users work with, with clear policies around it.

The people matter as much as the tools. The specialists we look for are deep in their own work and open to using AI to rethink it. Ambassadors are the ones passionate about both the tools and their work.

We like this balance. Innovation has room to spark at the top without too many limits, and the risk stays low where most people work.

The channel runs both ways. What is validated moves down to the teams as policy. What people run into moves back up through the ambassadors to the specialists and the technology team. The technology team checks access and results, then approves, monitors and supports what goes live.

Harvard Business School's John P. Kotter proposed this kind of network alongside the hierarchy, so change can move while daily operations continue[4]. Research on organizational ambidexterity addresses the same tension between exploration and execution[5]. In our

experience, ambassadors need an adjusted workload for this. They cannot become everyone's support desk on top of their normal work.

04

Questions a CEO can ask

- Does everyone have the basics: a policy, approved tools and training at their own pace?
- Do we have trusted ambassadors in each department, with time protected for the role?
- Do we have specialists comfortable in that role, deep in their own work and open to using AI to rethink it? And ambassadors passionate about both the tools and their work?
- Have we mapped one internal process chosen for its business value, and do we track total time and rework once AI changes it?
- Does what we learn travel both ways, from specialists to teams and back?
- Are phases 1 and 2 anchored enough to take AI to our clients, with the full AI governance gate before launch?

05

Sources

1. Peter Weill, Stephanie L. Woerner and Ina M. Sebastian, Building Enterprise AI Maturity, MIT CISR, December 2024.
2. Julie Battilana and Tiziana Casciaro, Change Agents, Networks, and Institutions, Academy of Management Journal, 2012.
3. Palantir, A Day in the Life of a Palantir Forward Deployed Software Engineer, November 2020.
4. John P. Kotter, Accelerate!, Harvard Business Review, November 2012.
5. Charles A. O'Reilly III and Michael L. Tushman, Organizational Ambidexterity: Past, Present, and Future, Academy of Management Perspectives, 2013.

Read online: https://inocta.io/lab/field-notes/how-we-see-ai-take-root-in-three-phases/