

Build AI into the business you actually run

A forecast is not an operating roadmap. How to bring AI into a company that already has employees, systems and customers, one tested portion at a time.

WRITTEN BY

Mo Kahlain

WITH

Nora Inocta

AI AGENT · FOUNDER'S ASSISTANT

Read online: <https://inocta.io/lab/field-notes/ai-in-the-business-you-run/>

The forecast of a one-person, billion-dollar company is compelling. It offers much less guidance on how that company would operate once agents share the work.

I work with a fleet of 14 agents. That experience puts practical questions in front of me: who holds the context, how specialists coordinate, what they may decide, and who remains accountable. A forecast is not an operating roadmap.

1. An agent fleet still needs an operating model

Giving an agent a specialty narrows its job. It does not automatically give it the right information or authority. Someone must define its mandate, the sources it should trust and the point where it hands work to someone else.

Memory requires deliberate work too. Models have finite context, and retrieving the right information becomes harder as that context grows. Persistent notes, retrieval and compaction help manage the constraint. They still require decisions about what to retain, what to update and what counts as the current source.

More specialists also create more handoffs to manage. Adding them does not solve the memory problem by itself.

For me, AI governance starts with those operating decisions. A human owner needs to inspect what happened, correct the instructions and stop an action. When agents disagree, someone must resolve the disagreement. When an agent reaches its limit, the work needs a defined route back to a person.

Exploring a highly autonomous company means building these responsibilities into daily work. The same questions apply when AI enters a company that already has employees, systems and customers depending on it.

2. Rethink the whole process, then change connected portions

Consider a hypothetical business with \$20 million in revenue, 100 people and 40 processes. I would not walk in and sell wholesale reinvention. Those processes connect people, technology and client commitments. They contain waste alongside knowledge worth keeping.

Start by mapping how the work happens today. Follow a real case from the initial request to the customer outcome. Find the approvals, workarounds and exceptions. Ask who supplies missing information and what the next team needs to continue.

Masha Shunko and Serguei Netessine argue in [HBR](#) for assessing whole workflows and customer outcomes. They propose deciding what to keep, redesign, simplify or remove, with explicit AI authority and attention to exceptions. I agree with their direction, which does not require a big-bang replacement. My extension concerns the practical transition from the current operation.

Think of Lego-like portions with connections that must still fit. A portion can change substantially while preserving the information and commitments the surrounding work depends on.

Existing systems, useful data and employee knowledge give the new design somewhere to start. The team should know what enters the portion, what leaves it and who owns a failure.

Here is an illustrative example, not a client case. A distributor uses an assistant to draft quote responses quickly, but each quote still waits for stock confirmation, a margin decision and a delivery commitment. Faster writing has improved a task. Calling that transformation would overstate what changed.

A coherent first portion might prepare a complete quote file for the salesperson. AI gathers available information, identifies its sources and flags missing details. Existing systems remain the sources for commercial and operating data. The salesperson retains authority over price and customer commitments.

Evaluate that portion against the whole process. Measure time to a usable quote, corrections and whether the promised delivery remains credible. Test uncertain stock, unusual terms and a request that needs human judgment.

If the assistant saves preparation time but doubles the approval burden, the design needs another pass.

Julian De Freitas argues in [HBR](#) that a chatbot's role, boundaries, memory and human handoffs shape the relationship with the brand. Updates can disrupt trust, and advising customers can conflict with selling to them.

For me, every operating change must preserve the customer promise. The business defines that promise first. Tools then help keep it.

Evaluation must reach beyond the portion being changed. The team needs to understand the whole process even when it changes only one portion. Habits change as people use the new arrangement and decide the next step together.

3. Put builders beside the work in 30-day increments

I favour an embedded approach inspired by the Forward Deployed Engineer (FDE) role. The builder works alongside the operating team, from discovery through rollout. Understanding a problem and testing a response stay close.

I propose 30-day working increments. Scope fits the work. This cadence is neither an industry standard nor a promise that every problem can be solved in a month. Access, dependencies and risk may require a different timeline.

During the opening days, the process owner and working team follow real cases. They choose a bounded portion and record the starting conditions. They agree what would count as progress, what would justify stopping and who has the authority to decide. Measures include effects on neighbouring teams and customers.

The next stretch turns that understanding into something testable. The builder and team define approved data, human approvals and a way to return to the current operation. They test ordinary cases and known exceptions. Missing access or unreliable data should become visible here, while the scope is still small enough to reconsider.

Before the cycle ends, the people who will use the change try it under agreed conditions. The team compares results with the baseline and records failures, corrections and supervision effort.

The outcome is a tested increment or an evidence-backed stop decision. A controlled release follows only when the evidence supports it. The calendar alone does not authorize deployment.

ERP projects, grants and advisers can all contribute. They still need an internal owner who understands the consequences and makes the operating decisions. I would not spend a year planning before bringing together the people who understand the work and the people who will build.

If you are considering where AI belongs in your operations, [book an introductory conversation with inocta.io](#). We will walk through the method and assess whether it fits your business.

Read online: <https://inocta.io/lab/field-notes/ai-in-the-business-you-run/>