FIELD NOTE

November 19, 2026
5 min read

I Built a Team of AI Agents. Then I Became the Bottleneck.

Fourteen AI agents work in our operations, and the consequential steps still wait for one person. What HBR and new research say, and what we are changing.

WRITTEN BY

Mo Kahlain

Read online: https://inocta.io/lab/field-notes/i-became-the-bottleneck-of-my-own-ai-agents/

On November 3, I asked one of our AI agents to review a draft of this post with me. It came back with a careful review: six marks, and each one was a thing for me to check. The agent did good work. It also handed the work straight back to me.

That is my week in one moment. At inocta.io, fourteen AI agents work in our own operations. They write, build code, handle part of the bookkeeping and follow up with clients. In the 30 days to November 3, they recorded 1,042 task closures. A large share of what matters still waits for one person before it moves, and that person is me.

01

The loop has a name

In Harvard Business Review, Sangeet Paul Choudary calls it the "human-in-the-loop" fallacy. Many companies place people around the tasks a machine cannot do yet, as a safety layer on top of the automation. He shows why that layer gets weaker over time. People who monitor a reliable system grow complacent. A wrong AI answer sounds as fluent as a right one. Every correction becomes training data, so the human role keeps shrinking. His word for these checkpoints is "unreliable and self-liquidating".

Research that Anthropic published on October 27 adds a sharper finding. Fifty professional analysts worked with an AI agent that the researchers had given a lean on purpose. Across four analytical choices, the analysts went with the agent's lean 86% of the time, where chance is about 50%. Experience did not protect them, and their confidence looked the same either way. The study covers one task, in one field, with one model, and the lean was made strong on purpose. It still tells me something uncomfortable. My sign-off does not prove that I judged.

02

How we built the loop ourselves

Nobody planned to put me in every loop. We drew our rules by type of action. Anything that sends, publishes, deploys, spends money or changes a core rule needs my yes. Each rule made sense on its own. Together they point every consequential step at one desk.

The loop is also harder to see than it looks. On November 3, our own ticket system showed two items waiting for me. The real count was larger. Five code changes were waiting for me to merge them. Four had been reviewed along the way. None had a passing review on its final version. Several tickets said "waits on Mo" while the field that tracks who acts next named another agent.

The hardest part for me is that each agent follows its rules correctly. It finishes the work and checks it. Then it asks before the step it may not take. The friction comes from the design, so the change has to come from me.

03

What we changed, and what we are testing

The first change moved the check off my desk. Since October 26, a high-stakes task cannot close without a passed independent review or a written, counted reason to skip it. The reviewer is another AI model or a separate agent. Choudary asks for more: rules embedded in the work and verified continuously. This is a first step, and the first count is not flattering. Our record of reviews started on October 26. By November 3, 31 of the 358 closures since then carried a recorded review. Of the 38 closures under the new rule, 28 used the written skip. Now we can see the gap.

The second change is on my table now: a named approver who is not the CEO. Take an internal change with a passed review and green tests, one that touches no client and no money. The agent that owns it would merge it. I would see the list each day, and any merge can be undone. The passed review must cover the final version, so none of those five would have qualified.

The third is about earned trust. Each agent has a score built from its record. We decided that the score opens no door yet. For 30 days it will only write down what it would have allowed, next to what really happened. That log has not started. Then I decide on real cases, not on a number.

Some calls rest on my judgment, such as a price or advice to a client. For those we are planning one more test. The agent would give the strongest case for each option before its pick. Where my own view is the evidence, it would ask for my view first.

04

Point the freed time somewhere

None of this matters if the saved time disappears into more checking. IKEA shows another path. Between 2021 and 2023, Ingka Group's customer-service chatbot solved 3.2 million interactions and saved nearly €13 million. Ingka also reskilled 8,500 call-centre co-workers in remote interior design, digital retail sales, relationship building and complex problem-solving. By the end of its 2022 financial year, sales through remote meeting points reached €1.3 billion. The chatbot removed work. The decision about where people went next created the value.

That decision needs an owner. We Map how the work runs today, then Structure who decides what. Only then do we Deploy a focused application and Manage it with a named owner. The Structure step is where most of the human-in-the-loop load gets set, often before anyone notices.

05

Five questions for your next leadership meeting

Before you expand an AI initiative, I would put these on the table:

1. What customer problem can we now afford to solve?
2. Who owns that outcome and has the authority to improve it?
3. Which decisions may the agent make on its own, which need a named approver, and is that approver someone other than the CEO?
4. What evidence would make us stop, change or expand the process, and who updates the instructions when the same correction comes back?
5. Where will the freed-up time go, and how will we know it created value?

If your honest answer to the third question is "me" for most of the list, you have moved the work to your own desk. I know, because I did it.

Read online: https://inocta.io/lab/field-notes/i-became-the-bottleneck-of-my-own-ai-agents/