

BEYOND THE CO-PILOT

Stop Calling AI a Co-Pilot — It Is Becoming Part of the Operating Model

The uncomfortable 2024 truth: once AI coordinates work, role boundaries and accountability must change with it

Giovanni Leonardi

April 2024



Stop Calling AI a Co-Pilot — It Is Becoming Part of the Operating Model

Transformation · No. 313

Originally written April 2024 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The uncomfortable truth is that AI is becoming part of the operating model while most organisations are still managing it as software adoption.”

1. The Co-Pilot Metaphor Has Run Out of Road

“Co-pilot” was a useful way to introduce generative AI.

It reassured employees that the human remained in control. It framed the technology as assistance rather than replacement. It encouraged experimentation with drafting, summarising, analysis and search. It made a disruptive capability sound like a better tool inside an unchanged job.

That is precisely why we now need to stop using it.

An AI system that waits for a prompt, produces a draft and hands control back may resemble a co-pilot. An AI system that receives a goal, gathers information, chooses tools, sequences tasks, asks for approval and continues until an outcome is reached does not. It has entered the workflow as an actor.

The organisation can still insist that humans remain accountable. It should. But accountability is not preserved by metaphor. It is preserved by redesigning authority, work allocation, supervision and evidence.

The uncomfortable truth is that AI is becoming part of the operating model while most organisations are still managing it as software adoption.

2. We Are Automating Tasks and Preserving Fictional Jobs

Most implementation plans begin with roles. They ask how an assistant can make the project manager, analyst, adviser or engineer more productive. Each role keeps its boundaries. AI capability is added inside the box.

This is administratively convenient and operationally weak.

Work does not arrive in neat role-sized units. A service request may need classification, research, judgement, drafting, approval, communication and follow-up. Today these steps cross several people and systems. An agentic workflow can increasingly coordinate

part of that sequence.

When it does, the unit of design is no longer the individual job. It is the flow of decisions.

Consider a composite team of nine analysts processing 1,200 policy queries each month. An AI workflow classifies new requests, retrieves approved guidance, drafts a response and routes unusual cases to a specialist. Drafting time falls by 55 per cent. The programme declares a productivity success.

Within six weeks, the team manager faces a different reality.

Junior analysts receive fewer ordinary cases and more ambiguous ones. The two specialists become bottlenecks because escalation volume rises. Quality review takes longer because drafts are fluent enough to conceal small errors. The manager cannot tell whether a delay belongs to the model, the retrieval source, the analyst or the approval rule. The team produces more drafts but resolves only 18 per cent more cases.

The tool improved several tasks. The operating model did not absorb the new distribution of work.

This is the pattern we must confront. AI removes routine effort unevenly. It changes who sees which cases, where judgement accumulates, how people learn and what managers must supervise. Leaving the organisation chart untouched does not preserve stability. It hides the redesign until it appears as workload, quality or accountability failure.

3. A Colleague Without a Role Is a Control Failure

If we want AI to behave more like a colleague, we must answer questions we would never leave vague for a human colleague.

- What work is it assigned?
- What information may it access?
- Which decisions may it recommend, prepare or execute?
- Who reviews its exceptions?
- How is its work attributed and reconstructed?
- What happens when policy changes?
- Who can stop it?
- Which person owns the consequence of its action?

These are operating-model questions.

Yet many organisations answer them in prompts, technical settings or supplier contracts. They do not connect them to process authority. A prompt says the agent should escalate high-risk cases, but the business has not defined “high risk”. A tool restricts transaction value, but nobody owns the cumulative exposure across thousands of small actions. A human approves outputs, but the approval queue makes the promised benefit impossible.

The technology exposes ambiguity that human work previously absorbed through experience and informal relationships.

AI does not merely cross role boundaries; it reveals that many of those boundaries were being held together by unrecorded human judgement.

4. Stop Calling Every Human Touch “Oversight”

The accepted answer to uncertainty is “human in the loop”. It sounds responsible and usually means very little.

Which human? At what point? With what evidence? How much time? What authority? What happens when the queue grows? Is the person verifying a fact, accepting a risk or merely clicking approval because the output looks plausible?

A human touchpoint is not a control unless the person can detect the relevant failure and act on it.

We should distinguish:

- verification — checking evidence against a defined rule
- judgement — choosing among legitimate alternatives
- approval — accepting responsibility for an action
- exception handling — resolving cases outside the AI boundary
- supervision — observing patterns across many runs and changing the system

These are different jobs. Combining them under “oversight” creates an invisible role with unlimited demand.

As AI handles more ordinary work, human work becomes more concentrated in exceptions, conflict and consequence. That work is harder, not lighter. Staffing models based only on hours saved will systematically underfund it.

5. The Serious Objection: Do Not Anthropomorphise Software

There is a strong objection to the “colleague” language. AI has no intent, duty or moral agency. Calling it a colleague may encourage trust it has not earned and allow leaders to pretend that accountability can be delegated to a machine. It risks disguising automation as partnership.

That objection is right about the danger.

We should never confuse fluent behaviour with personhood, or technical autonomy with accountable agency. The organisation remains responsible. A machine cannot accept the consequence of a decision.

But rejecting the metaphor does not remove the operating-model shift. Software that initiates actions, coordinates steps and shapes what work humans see is not a passive tool in practice. It changes the system whether we call it an agent, assistant or workflow.

The answer is not to humanise AI. It is to operationalise its role with greater precision than we often apply to human work.

6. Design the Human-AI Unit, Not the AI Feature

We need a new design discipline.

Start with the outcome and map the full flow of decisions. Assign each step to one of four states:

1. Human-owned. The person performs and remains responsible for the decision.
2. AI-assisted. AI prepares evidence or options; the person decides.
3. AI-executed under approval. AI prepares an action; an authorised person releases it.
4. AI-executed within bounds. AI acts inside enforced thresholds; a human owner supervises the service.

Then redesign the human work around what remains.

Who handles exceptions? How will junior people learn when ordinary cases disappear? Which specialist capacity becomes constrained? What evidence must accompany an AI recommendation? How does a manager review a hundred actions rather than one case? Which service measure prevents speed from concealing rework?

This is not a technology workstream. It is operating-model design.

7. Rewrite Roles Before Productivity Rewrites Them for Us

The profession is tempted to promise that AI will release people for “higher-value work”. That phrase is not a plan.

Higher-value work requires defined demand, capability, authority and measures. If those are absent, saved time becomes more volume, fewer roles or an overloaded approval queue. The organisation receives efficiency without transformation.

We must act differently:

- rewrite role descriptions around decisions and consequences, not inherited tasks
- fund exception handling and supervision as first-class work
- redesign learning pathways before routine cases disappear
- measure end-to-end resolution, not AI output volume
- make AI authority explicit, technically enforced and revocable
- assign a human owner to every operational outcome
- change team structures when work migrates across boundaries

The manager of a human-AI unit will need to understand process, data, model behaviour and human capability together. The specialist will need to teach the system as well as resolve cases. The junior practitioner will need structured exposure to judgement, not years of administrative repetition. These are new operating responsibilities, not side

effects.

8. The Call for Change

We should retire three comforting beliefs.

First, that adding AI to every role preserves the role. It preserves the box while changing the work inside and between boxes.

Second, that human accountability survives because a person remains somewhere in the process. Accountability survives only when authority, evidence and consequence are explicit.

Third, that productivity can be captured before the operating model changes. It cannot. The operating model determines whether saved effort becomes value or merely moves the bottleneck.

The move from co-pilot to colleague is not a claim that AI has become human. It is a warning that AI is beginning to participate in work with enough initiative to alter how humans coordinate.

We can continue deploying assistants and allow those changes to emerge as surprises. Or we can design the human-AI unit deliberately: bounded machine action, named human accountability, visible evidence and roles built around the judgement that remains.

Do not wait for the organisation chart to become obviously wrong.

The operating model is changing now. Our responsibility is to admit it—and design it before the technology does so by default.