AN ONLINE GUIDE · FIRST EDITION, NOVEMBER 2026

The One-Workforce Method

Managing a human-AI workforce

AI has become a new workforce category. From here on it is one workforce, and you cannot manage half of it and expect the whole to deliver.

Wawi Kantz

oneworkforcemethod.com

© 2026 Wawi Kantz e.U. All rights reserved.
The One-Workforce Method: Managing a human-AI workforce
First edition, November 2026
oneworkforcemethod.com
Wawi Kantz e.U., Lutzmannsburg, Austria

How to cite: Kantz, W. (2026). *The One-Workforce Method: Managing a human-AI workforce*. First edition. oneworkforcemethod.com

This guide and its templates are meant to be used. Apply them, adapt them to how your organisation works and share them with your team. If you want to build them into something you charge for, talk to me first. Wherever they travel, please keep the credit: The One-Workforce Method by Wawi Kantz, oneworkforcemethod.com.

Contents

Why I wrote this 4

PART I **The Claim**

1 Something has changed 9

2 Not a tool and not a person 14

3 The unit of management 19

4 The narrative has it backwards 25

5 Supervising and managing 30

PART II **The System**

6 Hired like a worker: the job description 38

7 Onboarding: the AI worker file 44

8 Performance review 50

9 The structured goodbye 52

PART III **The Roles**

10 The manager of an AI worker 55

11 The leader of a human-AI team 58

12 The enabler 62

13 The context owner 67

PART IV **Making It Work**

14 The upside and the hard questions 70

15 Starting small 73

16 Closing remarks 75

PREFACE

Why I wrote this

AI was never my plan.

Twenty years ago I started out in communications, marketing, and strategy. I loved my job, which kept handing me new opportunities, until the opportunity became a stretch. Half-way through my tenure at one of the largest pharmaceutical companies in the world, I stepped into AI with a strong business background ... and zero technical foundation. What followed was years of trial, error, and stubbornness. I taught myself to clean messy data, trained models that predicted everything from financials to customer preferences to disease risk, and led AI projects that had to solve real business problems. I helped shape how AI was governed across the company, and I built and led a global AI champion network of hundreds of colleagues supporting each other's adoption journeys.

I saw AI initiatives from every side: inside the project team, from leadership, through governance, and from the business that was supposed to benefit from the outputs. Something bothered me all along that I could not pinpoint. Despite doing everything the state of the art asked for (programmes, steering committees, user research, agile development, change management, you name it), I still felt like I was acting from an ivory tower at times.

That feeling became real one afternoon, when I was standing in an office kitchen and talking to a few of the people who were supposed to use an AI solution we had been working on for quite some time. We had done the work: bug fixes, the features people had asked for, training, office hours, adoption tracking. Then, between the coffee machine and the sink, I heard this:

In theory, of course it gets us to the draft much faster than I could. But in the end I'm responsible, and re-tracing what the thing wrote is as much effort as writing it myself. Feels a bit like hiring a new colleague who's supposed to take work off your plate, but then only causes more work because that newbie still has to learn a lot.

I chose my job for a reason. Now AI is doing the fun bit and I'm only nodding off. Not sure that's what I signed up for.

Do I even want to save hours of my work … would be kind of stupid to eliminate my own role, wouldn't it?

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

None of the people I talked to had been on the project team, and none had been in the user research. The research was good; it was simply pointed at the solution, and the ripple the new AI sent through people's work was not what we had initially gone looking for. The pattern behind it was one I had been part of myself, and it is the one that still dominates how AI gets launched today: we run it as a technology project.

AI is technology, and the decisions about data pipelines, integrations, permissions, and monitoring matter. I have spent years on them. But at the same time, we expect people to work collaboratively with these systems. I don't mean the way we work with email or any other tool. I mean the way we work with other human beings: briefing them, delegating to them, trusting them with a first draft, questioning their output, giving feedback so the next round gets better. We bring something into our teams that people are supposed to treat like a colleague … and we run it like infrastructure. Looking after the infrastructure is necessary work, and nothing in this guide argues otherwise. What we skip is the human side. The questions we would never skip for a new team member (who is responsible for you, what does good work look like, how are you doing, what changes for the people working with you) we rarely think to ask.

The strange thing is that we already know how to do this. Role clarity, onboarding, feedback, performance reviews, a structured goodbye when something no longer serves its purpose: for more than a hundred years organisations have been learning how to manage work and the people who do it. We have the discipline. We simply don't think of applying it to AI, because we keep telling ourselves AI is only a tool.

Getting this wrong rarely looks like failure. Nothing breaks. The system runs, the output is technically fine, and somewhere between the pilot and the second quarter it stops being used by anyone except the people who built it. Whether it is still doing good work is a question nobody was ever made responsible for asking. The team asked to work with it does what sensible people do: they check everything twice, or they quietly go back to doing the work themselves.

Then the verdict comes back that AI didn't deliver, and the answer is a better system: a more capable model, maybe a cleaner integration, and yet another round of budget. I've watched that answer several times. It produces a more capable system that disappoints in exactly the same way. Because none of it was ever a technology problem.

This guide sets out the One-Workforce Method. It takes what we already know about managing work and the people who do it, and tailors it to a workforce that now includes AI workers. AI has become a new category of worker, and from here on it is one workforce, people and AI workers together. You cannot manage half of it and expect the whole to deliver. And you cannot simply add new workers and not consider the ripple effect on the human workforce.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

A word on who this is for. If AI is part of how your work gets done, and you've wondered what it takes to make that collaboration actually work, for the results and for the people delivering them, this guide is for you. No technical background is required, and it applies whether you lead a large team or have no human reports at all.

Same principles. Different category of worker. Same but different.

That lens carries every chapter: here is where AI is the same as what you know, and here is exactly where it is different. It is also why we keep acting oddly around AI, sometimes reckless, sometimes too passive, forgetting the skills that got us here in the first place.

Across twenty years I've worked in different companies and departments, under many managers, and I've met every level of maturity these people processes come in. I've sat in organisations where every process existed on paper, and my development conversation was an admin task my manager clicked through between two meetings. And I've worked in places with barely any formal HR process at all, where a manager simply decided that my growth was a personal mission … and I grew more there than anywhere else. So yes: organisations own the discipline. But it is unevenly built, and even where it's built, unevenly meant. What made the difference in my career was never the process alone. It was whether someone behind it meant it. Which is exactly why I am going to spend this guide on process anyway. Conviction is fragile: it leaves when the manager leaves, and it rarely reaches the people two teams over. Process is how you give it somewhere to live.

Everything I'll ask you to build for your AI workers (role clarity, a real onboarding, honest reviews, a structured goodbye) is exactly what many of our people have never reliably received. If you catch yourself thinking *"we don't even do this properly for humans"* … good. That is the second opportunity this guide offers. Build the discipline once, and run it for both sides of your workforce.

Why write this now, and why like this?

For a few years my thinking on managing a human-AI workforce lived in many places at once: LinkedIn posts, a white paper, keynotes, client engagements, method documents, conversations. Piece by piece it held up; as a body of work it was scattered. So this is its home. I've written it like a book because a book forces coherence: every chapter has to carry its weight, and the seams between ideas have to actually hold. But I haven't waited for a publisher, and that's deliberate. The chapters live openly on oneworkforcemethod.com, each one readable on its own, and the whole guide can be downloaded as a printable file. Where I draw on research, vendor studies, or industry reports, the chapters say so.

I wrote this guide the way I work: together with AI. Drafts were sparred with, challenged, and sharpened by my artificial thought-challengers. On my own, I could not have drafted this at this pace. That speed has a backstory, though. Long before the first chapter existed, I

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

invested heavily in building up the context that lets my AI understand my concepts, follow how I think, and write the way I write: my frameworks, my positions, my voice, my rules for what good looks like, deliberately curated and continuously maintained. In other words, I did for my AI exactly what this guide will ask you to do for yours. I gave it a proper job description, onboarded it into my work, and I review everything it produces. The thinking, the experience, the positions, and every final word are mine. And if you find em dashes, they are mine too. I fell in love with them before AI slop became a thing, because they make a long thought easier to follow. And so is the accountability for what's in here, which, as you're about to read, is rather the point.

Wawi Kantz
Austria, 2026

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART I

The Claim

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART I · CHAPTER 1

Something has changed

“How many AI solutions do we actually have running right now? And what are they delivering?”

A leader asks it almost casually, in a meeting or on the way out the door. I know what happens next, because I’ve been the person it happened to. Emails fly across departments, someone digs out the inventory (if one exists), and it turns out to hold half the systems and almost none of the answers. Even with a starting point, you usually have to collect much of the information by hand. The second half of the question, what they are delivering, often has to be retrofitted to a metric that was never set up in the first place.

That was the situation before generative AI, and it was hard even then. But back then we had something working in our favour that I only learned to appreciate once it was gone: friction.

Until a few years ago, an AI solution couldn’t just appear. It took a budget someone had to defend, a team with data science or engineering expertise, or a vendor brought in through procurement. There were reviews and steering meetings. We complained about all of it, and I did my share. But it left a trail. Somewhere there was a contract, a charter, a deck, a name. An overview was hard to get, but in principle it could be reconstructed.

Today the trail is often never created at all. Any reasonably capable colleague with a GenAI tool can build their own solution: an agent that drafts the newsletter, a workflow that pre-sorts complaints. No coding needed, and usually no knowledge of technology lifecycle management, and no compliance beyond the mandatory how-to-use-AI-responsibly training.

If you have ever been responsible for AI adoption, you know the tension this creates. The same role that asks you to fuel experimentation, by celebrating the colleague who built her own agent, is also the one supposed to know what all that enthusiasm produces. Push the controls too hard and you strangle the initiative you spent years encouraging. Hold back and things multiply that nobody can see. I won’t pretend I ever walked that line perfectly. If you have found the holy grail of getting this balance right … please, tell me.

Think about what that means in workforce terms. It’s like giving every employee the green light to hire whoever they want. There is no contract, no manager, and nobody tells HR. The new colleague starts real work the same afternoon: drafting customer-facing text, taking steps on its own. And these are exactly the “colleagues” the rest of our people are then asked to collaborate with.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

The instinctive fix is a better inventory: a registry, a control plane that captures everything. That is useful, and an entire market is being built around it. Regulation points the same way: in Europe, the AI Act is beginning to ask organisations for this kind of oversight. But a law can oblige you to watch your AI … it can't manage it for you. And a registry answers only one question: what exists. Imagine for a moment that the inventory were perfect, with every system and its owner listed and no blank fields. The second half of the leader's question would still be open.

Most AI that came out of a proper project does have an owner, usually a product or business owner. A good owner watches whether the output holds up. What tends to sit outside their reach is the other half: what the system changed for the people working around it. Whose job got rearranged, who is now checking everything twice, and what that person should be growing into instead.

For everything that was never a project, the picture is worse. The agent a colleague built last month has no stated owner. Whoever built it is the owner by default, usually without being told and without anyone deciding what that means.

What is it delivering, and is it any good? Who manages it? Those questions come from a different discipline entirely. They are the questions we ask about our people.

In a well-run organisation, a new hire with no job description and no manager would be noticed on day one. With AI, these things can go undetected forever. That contrast is what this guide is about. Everyone can hire now. What nobody handed out with that permission is any expectation of what comes after the hiring. There are governance rules about what you may use and which box to tick. About the work itself, there is nothing. So these AI workers stay unmanaged until somebody decides otherwise.

That decision may well be yours. If you built something that does real work for you or your team, and you can say what value it brings, the managing role is yours to step into. If you built it for somebody else to use, the conversation you owe them is about who takes that role day to day.

To see why "managing" is the right word, we need to look at what changed.

From answering to doing

For most of my career in AI, we saw the systems we built as tools. A predictive model suggested which message would resonate with which customer. You gave it a customer segment and it gave you a recommendation. Managing it meant monitoring accuracy and drift. IT-asset thinking fit, because they behaved like assets.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Then large language models arrived, and software could suddenly be briefed in plain language. The deeper shift followed: agentic AI, systems that research a topic, draft the deliverable and come back with finished work. Something that decides and acts is very hard to keep thinking of as a tool. A hammer doesn't decide which nail comes next.

The quieter shift matters most. AI capabilities became modular: the *skill*, the ability to do a kind of work, can now be separated from the *context* it works in. The same content-creation capability can serve Client X and Client Y, and you swap the brief underneath. Improve the skill, or correct something in shared context, and every worker running on it gets better at once.

That mirrors how human expertise works. A good marketer doesn't reinvent her craft for every client. She brings the capability, you provide the context, and the better the briefing, the better the work. Once skill and context come apart, the quality of what AI produces depends on what a person tells it and what they expect back. That is management. The moment AI started behaving like that, composable, briefable and able to act, it stopped being tool-shaped and became worker-shaped.

Worker-shaped, and everywhere at once

"AI is like an employee" is a line that gets thrown around loosely, and loose versions of it cause real damage. A later chapter covers where the comparison must stop. The observation here is narrower: the shape of the work AI does has changed, task by task, inside real jobs.

The business world has started saying it out loud. The biggest AI companies sell "agents" you "onboard," and enterprise platforms offer role definitions and lifecycle management for AI. Ethan Mollick put his finger on it: the decisive skill of this era is management, which means knowing what you want and checking what comes back.[1] When I first landed on that idea myself, it sounded almost too unglamorous to be the answer. I've since watched it hold up across every engagement I've run.

The claim

Everything that follows stands on one claim:

AI has become a new workforce category. From here on, people and AI form one workforce, and you cannot manage half of it and expect the whole to deliver.

Workforce category means AI takes on work that used to belong to people, and needs what workers need: role clarity, feedback, and someone responsible for it. *New* means it is nothing like a person: it has no stake of its own and no accountability. *One workforce*

means the halves can no longer be separated in practice. The work passes between them every day, so managing only the people leaves you managing one contributor to a shared result. The discipline has to run on both sides, from the job description at the start to the structured goodbye at the end, on the same terms you already use for everyone else.

The obvious objection is “Isn’t that overhead? The whole point of AI is speed.” If only it were that easy … In 2024, MIT researchers pooled 106 experiments in which humans and AI worked together. On average, the combinations performed *worse* than the better of the two working alone.[2] The losses were concentrated in decision-making tasks, and the gains showed up in content creation. So the effect depends on the kind of work, and the combinations that worked were the ones designed for the task at hand.

That is the argument in one result. The value of human-AI collaboration comes from decisions about roles, allocation, and oversight. There is a name for making those decisions deliberately, and it has existed for a century: management. When AI disappoints, and enterprise surveys are full of abandoned pilots right now,[3] I no longer ask what is wrong with the technology. I ask who is managing it, and against what expectations. In most organisations I meet, the honest answer is still nobody.

Same but different

The management discipline AI needs is one organisations already own. We know how to write role descriptions, review performance, and retire what no longer works. That is the *same*.

The *different* matters just as much, and ignoring it is how companies get into trouble. AI is not a person, and never will be. It doesn’t consent and it doesn’t carry responsibility; a human is always accountable for what it does. The craft lies in holding both halves at once: applying the discipline we know, while never confusing the new worker with a colleague.

Same but different. If you keep only one phrase from this guide, that’s the one I’d choose for you.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Take this with you: a 60-second check for your own organisation. Can you say how many AI systems are doing real work for you right now? For each one, can you name the person who manages it, and does that person know they do? Could you state, in one sentence each, what “performing well” means for them? When were they last reviewed? Has any AI system in your organisation ever been deliberately retired? If you answered “no” or “I don’t know” more than twice … you’re exactly who I wrote this for.

Let’s start with the question hiding inside all of this. What, exactly, is a worker?

NOTES

1. Ethan Mollick, “Management as AI Superpower,” *One Useful Thing*, 27 January 2026.
2. Michelle Vaccaro, Abdullah Almaatouq & Thomas Malone, “When combinations of humans and AI are useful: A systematic review and meta-analysis,” *Nature Human Behaviour* 8 (2024), 2293–2303.
3. S&P Global Market Intelligence (451 Research), *Generative AI shows rapid growth but yields mixed results*, October 2025, Figure 2. Among respondents whose organisations had AI or machine learning in production or at proof-of-concept stage, the share reporting that most of their AI proof-of-concept projects were abandoned before production rose from 17% in Q4 2023 to 42% in Q4 2024.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART I · CHAPTER 2

Not a tool and not a person

In July 2024, an HR software company called Lattice announced something that sounded, at first, almost reasonable: it would give AI agents official employee records and a place in the org chart, right next to the humans.

The backlash took about three days to force a full retreat. HR professionals were appalled, and the CEO ended up explaining that it had all been “a big misunderstanding.”[1] The company had, everyone agreed, crossed a line.

What fascinates me is what came next. Within two years, the biggest names in enterprise software were selling … almost exactly that. Workday and OpenAI both sold agents as workers to be onboarded, measured and managed. Microsoft told the world to govern agent identities “the same way you govern your employees.”[2] The idea that got Lattice pilloried became what the big vendors now sell, with barely an apology to the timeline.

So which is it? Was Lattice wrong, or early?

My answer: both, and the difference between those two words is precisely what this chapter is about. Because the storm around Lattice was never really about software features. It was about a category question we still hadn’t answered — what *is* a worker, and what would it mean to let something artificial into that category? Get the answer sloppy, and you either insult your humans or mismanage your AI. Get it precise, and a lot of confusion simply dissolves.

So let’s do something slightly unusual: find out what a worker actually is.

Two boxes, and neither one fits

When something new arrives in an organisation, we reach for the nearest box. AI comes with two on offer, and I have fallen into both of them personally, which is how I know what each one costs.

The tool box

Start with mine, because it’s the more embarrassing.

A while ago I built myself a research agent. Every morning it went out, scanned what had been published in my field overnight, and produced a report. It worked. The thing I had asked for arrived, on time, every day, for weeks.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

I never read them. A skim at best, usually not even that — a daily report on top of client work is not a thing a person gets to. The agent was fine. The reports were probably fine. What had never happened was anyone deciding what the work was for, how much of it was worth having, or what a good version would look like. “Daily” was a setting I picked in about four seconds on the day I built it.

What fixed it was not a better model or a cleaner prompt. I gave it a purpose narrow enough to be useful, moved it to every second week, and told it what I actually wanted to know. The reports got shorter and sharper. I read them now.

That is the tool box. You deploy it, you check that it runs, and nobody is responsible for asking whether the work is any good or whether it needed doing at all. Now multiply that by an organisation where anyone can build one of these in an afternoon.

The colleague box

The other box is subtler, and I experienced the problem it brings first hand while writing this very chapter. The AI and I had been working on a passage for quite some time, but at some point I realized it did not add any value and decided to cut it. And guess what? I felt bad breaking it to the AI, because we had been co-creating and finessing this now meaningless passage over a week.

The cut was right. But that feeling is the whole problem in miniature, and the research says I am nowhere near alone in it. In May 2026, BCG researchers published a piece in *Harvard Business Review* titled “Why You Shouldn’t Treat AI Agents Like Employees.” Where organisations anthropomorphised AI agents — treated them as colleagues — individual accountability dropped and review quality fell.[3]

You do not defend against that with willpower or with knowing better. I had both, and still felt the tug. If we start personifying AI, we might forget that we are dealing with a new kind of worker, not a human, one that cannot be accountable for what it produces.

The third thing

Neither box fits, and being precise about why is what makes the rest of this guide work.

Put the work in one column and the person in another. Everything that describes **the work** — a brief, and a standard to meet — sits in the first column, and AI matches it, often to a degree no human does. Everything that describes **the person doing it** — a life the work is part of, and a stake in its outcome — sits in the second, and AI misses it entirely, and always will.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

That asymmetry is the precise meaning of the phrase this guide keeps coming back to. AI is functionally a worker and categorically not a person. It's the same, on exactly the dimensions that management operates on. It's different, on exactly the dimensions that rights, protections, and personhood operate on.

And that is why "a new workforce category" is the only framing I have found that holds. The work needs managing, which rules out the tool box. The second column stays empty, which rules out the colleague box. What is left is a third thing — one nobody has had to manage before, and one that needs managing all the same.

Which is also why that BCG research, which at first glance seems to challenge my thesis, reads to me as fitting it more than undermining it. Look carefully at what failed there. What failed is treating AI as a *person*: the trust and the collegial courtesy, everything from the second column. Nothing in it indicts the first column. Job descriptions and performance reviews are the management disciplines, and the same body of research keeps finding that *they* are what make human-AI work succeed. The line the evidence draws is the line this chapter has been drawing all along:

Manage AI like a worker. Never treat it as a person.

The Lattice story makes perfect sense through that lens. Putting agents in the org chart with goals and reviews — management discipline, the first column, and the industry adopted it within two years because the need is real. Announcing it in language that put AI *next to the humans*, as fellow employees — a personhood signal, the second column, and people recoiled correctly. Lattice was early about the discipline and wrong about the category, in a single press release.

The line that must hold

Could the second column ever fill up? Could AI become a person in the eyes of the law?

It has been proposed. In 2017, the European Parliament floated the idea of granting the most sophisticated autonomous robots the status of "electronic persons", primarily as a means of addressing liability when autonomous systems cause damage. In a 2018 open letter, hundreds of experts from AI, robotics, law and other disciplines pushed back, arguing, among other things, that legal personhood modelled on humans would bring human rights with it.[4]

I think they were right, and for the working world I would go further: the empty column is a feature worth defending. Everything this guide builds works *because* it stays empty. You can retire an AI worker without a severance conversation, and rewrite its role overnight. The management discipline transfers precisely because the moral weight doesn't.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Which leaves one thing that can never be handed over with the work. An AI cannot be accountable for what it produces, because accountability needs something that can carry a consequence, and nothing about an AI can. So the accountability sits with a human. Always. Someone is answerable for every piece of work your AI workforce delivers, and a central job of management now is making sure that person is named and able to carry what they have been handed.

Humans are always accountable. AI never is. Building an organisation where that is operationally true — and bearable for the people in it — is where we are headed.

First, though, a question almost nobody asks, although everything after it depends on the answer: what counts as *one* AI worker? The chatbot? The model? The whole system? It sounds like a technical question. Getting it wrong quietly breaks everything you build on top.

Take this with you: Pick one AI that does real work for you and draw two columns. On the left, what it is being asked to deliver — the work itself, and what a good version of it looks like. On the right, what you would expect from a person holding that role: judgement about whether the work was worth doing at all, and responsibility when it goes wrong.

Mine, filled in, looks like this. Left column: a short brief every second week on what has moved in human-AI workforce management, with sources, no longer than a page — and a good version is one I actually finish reading. Right column: deciding whether the topic still matters to my business, and carrying the consequences if I repeat something wrong in front of a client. I also need the standing to tell me the brief is pointed at the wrong question. None of those can be handed over. They are mine — and knowing that is what makes the left column something I can delegate without worrying about it.

Now mark everything in your right column that your AI can never supply. The left column is your management job, in full. Every line you marked on the right has to be carried by a human — and it is worth being able to say which human.

NOTES

1. Lattice announcement and reversal: *Fortune*, 12 July 2024; SHRM, “Lattice Scraps Plans to Treat AI Bots as Employees After Backlash”; CEO interview, *HR Brew*, 12 August 2024.
2. Workday Agent System of Record (generally available February 2026); OpenAI Frontier launch coverage, TechCrunch, 5 February 2026; Microsoft Entra blog, “Govern AI agent identities and access the same way you govern your employees” (2025). Vendor positioning throughout — treated here as evidence of the category shift, on which more (including what these platforms don’t do) later in the guide.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

3. Kropp, Bedard, Wiles, Hsu & Krayer, “Research: Why You Shouldn’t Treat AI Agents Like Employees,” *Harvard Business Review*, 6 May 2026; James O’Donnell, “AI agents are not your ‘coworkers’,” *MIT Technology Review*, 29 June 2026.
4. European Parliament resolution of 16 February 2017 (2015/2103(INL)), para. 59(f); Open Letter to the European Commission on Artificial Intelligence and Robotics (2018), robotics-openletter.eu.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART I · CHAPTER 3

The unit of management

For a good part of my career, I was the alien in the room. I came into the world of data and AI from communications, marketing, and strategy — and when you enter an IT organisation with that background, you notice things the natives have stopped noticing. What I could never get past: how much energy technical teams pour into fighting about what things are *called*.

This started long before AI. I've sat through entire meetings where smart, expensive people argued about whether the thing we were building was a data catalogue, a data dictionary, or a data inventory — while the question of what anyone would actually *use* it for rarely made it onto the agenda. That question was supposed to be answered by the right definition. And AI made it worse. Model or assistant? Assistant or agent? Agent or skill, workflow, multi-agent system? I've watched teams lose time and energy to these debates, with every vendor and every new framework redrawing the definitions just as everyone had settled in.

Somewhere along the way it became my broken-record sentence — the one former colleagues could probably still recite in my voice: can we PLEASE stop fighting about terminology and just say what this thing is supposed to do — and why we even need it?

Back then, I meant it as a plea for meeting sanity. It took me years to realise I had accidentally been giving the answer to one of the most consequential questions in AI workforce management — a question almost nobody asks. In all the engagements and conversations, I hardly ever hear anyone wonder what exactly the unit is that we're supposed to be managing. It sometimes surfaces in different clothes — during an inventory attempt, when someone asks at which level an AI system should be documented. When it comes to the majority of use cases we find nowadays, the ones created through citizen development with hardly any formal project around them, the question of the unit mostly isn't raised at all. What exists instead is an undifferentiated pile of chats, custom GPTs, agents and automations, spread across whatever platforms your organisation licensed, that nobody has ever had to draw a boundary around.

The platforms have caught up somewhat: most of the common ones will now list every agent, who built it, how often it runs, what it costs. Useful, and still not management. The narrower point here: a complete list doesn't tell you where one worker ends and the next begins.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

And that is exactly part of the problem. You can't write a job description for undefined "stuff." You can't review its performance, and no one can be clearly accountable for it. So this chapter asks the question on purpose. What counts as *one* AI worker, and where does one end and the next begin?

That question decides everything. A job description has to be a job description *of something*. A performance review needs a something whose performance it reviews. A named human needs to know what, precisely, they are accountable for. Get the unit wrong and every practice in this guide still happens … it just happens to the wrong thing.

Without a defined unit, the same AI landscape produces wildly different numbers depending on who counts and how. Three, if you count the platforms people work in. Forty, if you count everything running inside them. Two hundred, if you count each team's deployment of the same capability separately. Every one of those counts can be defended. None of them tells you what the leader actually wanted to know.

Remember the fire drill that opened this guide — the leader asking how many AI solutions are running and what they deliver? Underneath that scene sits a harder problem: without a definition of the unit, the question isn't just difficult to answer. It isn't even well defined.

The tempting answer

The obvious move is to define the worker by the technology. One model, one worker. Or one agent, one worker. The vendors hand you the vocabulary ready-made — model, agent, skill, workflow, multi-agent system — and I understand the pull of using it, because at least it sounds precise.

I tried it that way myself, but there are two problems that come with this approach.

The first you already know from my broken record: the vocabulary won't sit still. What we called a model became an assistant, became an agent, became a fleet of agents with an orchestrator — and by the time you read this there will be a new fancy term for a new fancy development.

The second problem is worse: the technical units don't line up with anything you'd want to manage. One foundation model can power fifty completely different solutions across your organisation — is that one worker with fifty jobs? A single "agent" can be anything from a trivial email sorter to a system that handles your entire customer intake. The label tells you nothing about the work, the risk, or the responsibility. Counting agents is like counting your workforce by number of hands: technically a number … but from a management perspective, utterly useless.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

We don’t manage human workers by their anatomy either. We manage a *role* — and a role is something organisations construct deliberately.

The human lesson: jobs were always designed

Step back to how the human workforce is actually organised, because the answer has been sitting there all along.

Work doesn’t arrive from nature pre-cut into jobs. A “job” is a bundle of tasks that an organisation has chosen to group together, attach an outcome to, and hand to one person with a name, a manager, and a salary line. Different companies cut the same work differently; the same company re-cuts it every reorganisation. The job is a management construct — and that is its strength. It’s the level at which expectations, feedback, and accountability can actually attach to something.

A human job is typically bundled for reasons that have nothing to do with the outcome. A person has forty hours a week to fill, so you keep adding until the week is full: one person’s job could encompass protocol writing, booking business travel, creating presentations, and tracking budget. Nobody would hire someone to write protocols and do nothing else.

AI has no week to fill. An AI worker can be exactly as wide as its outcome and no wider. Which means the units you cut will often look too small beside a human role — and that is the point rather than a problem.

So when the unit question finally gets asked — what counts as one AI worker? — my answer is: the same thing that counts as one human worker. A deliberately constructed unit of work and accountability. You don’t find it in the architecture diagram. You decide it, the way organisations have always decided what a job is. But with AI, there are no working hours to account for in your decision.

You don’t have to decide blind, and you don’t have to decide once. Cutting an AI role takes experience, and it depends on how your human workforce is organised as much as on the AI. The scope will very likely change over time, which is fine as long as each change is a decision rather than a drift. Three questions to start from:

The three questions

One: does it deliver a distinct outcome someone would miss if it stopped? An AI worker earns its place through an outcome, and “distinct” is the operative word. If a newsletter stops going out, someone would notice. If customer enquiries were no longer triaged, customer service would feel the pain. And if protocols stopped being drafted, you would be frustrated doing it manually again.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

If nobody would notice your AI worker stopping, you're either looking at a fragment of something bigger … or at a worker whose retirement is overdue.

Two: can one named human be accountable to manage it day-to-day? That should be a single name, never a committee or "the AI team." Human workers aren't managed by committees either. And that human needs the knowledge and skills to oversee the AI worker's work and its outputs — plus the ability to judge whether it delivers what it is supposed to, when it needs refinement or additional tasks, when there is trouble collaborating with its AI or human co-workers, and when it is time to retire it.

Three: how relevant is a consistent way of working? This helps you decide at what point it makes sense to split one worker into many. Take a protocol writer as an example. If your goal is a consistent way of writing protocols across your organisation, then it makes a lot of sense to have one AI protocol writer, managed by one human who decides what good looks like and makes sure the worker collaborates well with all the teams. If the way protocols are written is genuinely specific to each team, it makes more sense to treat each team's protocol writer as its own worker, managed by someone from that team. None of that stops you reusing the underlying technical components — transcription capabilities, skills, and the rest.

Once all three land, you have the level to write the job description at, name the human at, and review at.

And notice what the questions never ask: what the technology is. If a single humble "skill" — one well-briefed capability doing one job — passes all three, it's a worker, no matter how small it looks next to the shiny multi-agent platform. If it takes a five-agent system with an orchestrator to produce one distinct outcome under one accountable human, then the *system* is the worker and the five agents are its anatomy. You don't performance-review your colleague's left hand.

The technology is the anatomy. The worker is the unit of accountability.

Running the test

Let me make that concrete with three cases I keep meeting.

The newsletter agent — the one a capable colleague built without asking anyone, and our recurring guest. Outcome: clearly yes, the newsletter. A named owner: nobody today, but there *could* be one, and the question asks only whether that is possible. Its own briefing documents and subscription seat make a consistent way of working easy to define. That unassuming little agent is a full worker, subject to the whole management lifecycle this guide describes … which should change how casually it came into the world.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

The customer-service system — intake agent, retrieval agent, drafting agent, escalation agent, orchestrator. Run the test on any single agent and it fails immediately: no distinct outcome anyone would miss on its own, and nobody who could sensibly be asked to manage it alone. Run it on the system and everything snaps into place: one outcome (resolved customer inquiries), one accountable owner, one way of working, one cost line. One worker, five organs.

The foundation model — the big general-purpose model behind dozens of your solutions. Distinct outcome? It has none of its own; it has whatever the solutions built on it deliver. One accountable human for everything it touches? Impossible, and wrong to attempt. The model fails the test on every question — because it was never a worker. It's closer to what education is for humans: the general capability your workers draw on. Two humans with the same degree are still two different workers, doing different jobs in different contexts. Two AI solutions built on the same model are exactly the same case. The worker lives where capability meets a defined role — never in the capability alone.

The other cutting line

Your worker unit and the unit the law regulates are not always the same. The management unit is something you decide. The regulatory unit depends on where you operate and on how the system is provided and used. The EU AI Act, for example, defines an AI system by its characteristics, not by how many components it has or how many agents it runs. Other jurisdictions might draw the line elsewhere, so check the definition that applies to you.

What this unlocks

Go back one last time to the fire drill. "How many AI solutions do we have, and what are they delivering?" At the start of this guide, that question had no answer. Now watch it become answerable: count *workers* — units that pass the three questions — and suddenly there is a number that means something. Each worker on the list has an outcome you could state, a name beside it, a review you could schedule, a cost you could defend. The forty entries in the architecture diagram might collapse into nine workers; the two hundred "agents" might turn out to be twelve. The number gets smaller and the truth gets bigger.

That list — workers, not systems or their technical components — is the founding document of everything that follows: each one gets hired properly, onboarded, briefed through deliberately governed context, reviewed, and one day retired with intention.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

But a list of workers immediately raises the next question: what work do we actually give them — and what stays with your humans? If AI really is a new category of worker, there's no reason to believe today's division of work, which nobody ever designed, got that answer right. That's next.

Take this with you: Pick the messiest corner of your AI landscape — the one nobody can quite describe — and list the candidate units, whatever they're called technically. Run each through the three questions: distinct outcome someone would miss? One nameable human to manage it day to day? Its own way of working, its own cost? Merge the fragments, split the blobs, and see what number you land on. In my experience the count often shifts substantially in either direction … and the conversation it triggers about question two — the name — is worth the whole exercise.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART I · CHAPTER 4

The narrative has it backwards

Let’s take stock, because the first three chapters have quietly rearranged the furniture. AI is doing worker-shaped work. It passes every functional test of what a worker is, and fails every personhood test — manage it like a worker, never treat it as a person. And you now know how to define what *one* AI worker actually is. Which means your organisation has a combined workforce: human workers and AI workers, side by side, whether anyone planned it or not.

Sit with that for a moment, because a question follows from it — and I find it equally exciting and uncomfortable. If the workforce has changed, then the way work is currently divided between your humans and your AI is very unlikely to be the right one. Nobody designed that division. It grew. Tasks got automated where automation happened to be easy, or where a vendor happened to have a product. Citizen developers built what they personally needed. And the humans kept everything else — by default, because they were already doing it. The allocation of work in most organisations is a historical accident wearing the costume of a decision.

We would never accept that for a human team. When a new colleague joins, we re-divide the work. And AI is a stranger case than any new colleague: a category of worker with a completely different profile of strengths, different failure modes, and — remember the protocol writer — none of the bandwidth limits that shaped every org design decision we’ve ever made. When something like that enters the workforce, the honest response is to re-ask the composition question from scratch: given everything this work requires, who — human, AI, or both together — is best placed to do what?

If that question sounds familiar, it’s because the market has been shouting an answer at you for two years: AI first.

Any “first” implies a queue

I am a big believer in the power of AI. But I’ve always struggled with the term “AI first,” because it creates a competition nobody asked for. Two words, and you’ve told your whole organisation where humans rank. I know that’s rarely the intention behind the phrase … still, language is powerful, and your people hear the ranking long before they hear the nuance.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

The polite alternative doesn't rescue it. "AI augments humans" sounds kinder — and I think it shares the core mistake of the frame it's trying to correct. The augmentation frame quietly positions humans as the baseline and AI as the helpful addition at the margins, which keeps you comfortably in the old world: work defined by what humans have always done, AI assisting at the edges. Flip the words to "people first" and you're playing the same game with the queue reordered. Any "first" implies a queue. And both frames answer the question of *who* before anyone has asked what the work actually is — which means neither is an answer to the composition question at all. They're ways of avoiding it.

The question that eats the composition question

And beneath both frames, doing the actual steering, sits the question they collapse into in practice:

"Where can AI save us money?"

Sometimes it's spoken out loud, sometimes it hides behind more polished formulations, but in the end that's the question most AI conversations with a business leader boil down to — and it swallows the composition question before anyone can ask it. I understand where it comes from … budgets are real, margins are real, and the pressure to show returns on AI spending is very real, too.

The question deserves an answer, and I would never tell a business to ignore cost. What I've stopped believing is that this question can carry an AI transformation. Because look at what it quietly does: it starts with the technology and goes looking for a purpose, and it looks through one lens only — making the things you already do cheaper and faster. It takes the accidental division of work as given and optimises inside it. You polish the current game. You never ask whether it's still the right game to be playing.

And the damage goes beyond strategy. When the dominant narrative inside your organisation is "find what humans do and make it cheaper," your people hear the message precisely: the goal is to need less of them. That's where the "AI is coming for your job" anxiety gets its fuel — often not from the technology itself, but from the framing leadership chose. Adoption gets harder, trust gets nearly impossible, and the very people whose engagement you need start protecting themselves instead. There's a competitive angle too: if every organisation automates the same tasks to save the same costs, the savings get competed away and nobody gains a lasting edge. A whole economy optimising yesterday.

What the returns actually say

If the replacement narrative worked, it would at least show up in the numbers. It doesn't — and the gap is instructive.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

In its latest global survey of more than three thousand organisations across twenty-four countries, Deloitte found that two-thirds report productivity gains from AI. Only a fifth see revenue growth. And only about a third are doing what the researchers call transformative redesign of work — actually changing how the work is composed, rather than sprinkling AI on top of it.[1] What that tells me: most of the gain disappears somewhere between the activity and the P&L. Picture a team that saves four hours a week. Nobody decides what those hours are *for*, so they dissolve into the general busyness. Freed-up time is spare capacity, and spare capacity without redirection is exactly worth zero. The hours-saved slide looks great precisely because nobody asks the follow-up question.

McKinsey's research puts a finer point on where the blockage sits. In its 2025 report on AI in the workplace, just one percent of surveyed US executives described their organisation's generative AI rollout as mature, and the researchers' blunt conclusion was that the biggest barrier to scaling is leadership, not employees.[2]

My favourite detail from that study is the perception gap: C-suite leaders estimated that four percent of their employees were already using generative AI for more than 30 percent of their daily work. The employees' own reports? Thirteen percent, more than three times as many. The colleague quietly building their own agents and workflows in ChatGPT or Copilot is not hypothetical. The numbers show them too.

And the replacement narrative has one more problem: it misdescribes what AI actually does when it enters an organisation. AI takes over parts of jobs, and what is left behind is a role that needs recomposing — design and management work, precisely what the cheaper-and-faster narrative never budgets for.

Start with the work

So if every "first" is an avoidance, start somewhere else. Start with the work.

The sequence I use with clients is unglamorous and it changes everything. First the *why*: where is this organisation actually trying to go — the strategic goal, the game you want to be playing? Then the *what*: the work that needs to happen to get there, and the full set of capabilities that work requires — including the ones you've never had. Only then the *who*: for each capability the work needs, where does the strongest and most reliable version of it live? Sometimes the answer is clearly AI — processing and predicting at a scale and consistency no human can match. Sometimes it's clearly human — the judgment and relationships that work through ambiguity and politics. And often it's a combination, where the real design question begins: what's the human's role and what's the AI's role *within* this capability?

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

The word that carries the whole method is “honestly.” Honest mapping means resisting two defaults at once: the inertia of “a human has always done this, so a human should keep doing it,” and the hype of “AI can technically do this, so let’s automate it.” Both defaults feel safe in the moment. Neither asks what the work is best served by — and that question is the only one that matters.

What you find when you start with the work

The reframe earns its discomfort in practice. When I sit with teams and ask about the work itself, the same three patterns keep surfacing.

Dormant high-value skills. People have capabilities — strategic, creative, relational — buried under repetitive operational work. The day-to-day leaves no room for what they’re actually best at. AI can clear that ground. This is the efficiency play, it’s real, and notice: it’s the only one of the three patterns the money question can find, because it’s the only one visible in existing workflows.

Missing capabilities. Things the team simply cannot do: the analysis nobody can run, the expertise parked under “if we ever have budget.” Opportunities never pursued because they seemed out of reach. AI changes that equation entirely — this is the expansion play, and it’s where the transformative value lives.

Inconsistent skills. The same skill exists across the team, but quality varies wildly between people. Everyone has their own way of doing it, there are no shared standards, and when the best person leaves, their way of doing it walks out the door with them. AI can capture the best approach and make it the baseline — raising the floor so the whole team operates at a consistently high standard. This is the standardisation play, the quietest of the three and one of the most impactful, and its organisational form looks like this: one worker, one way of working, serving every team.

Replacement logic finds the first pattern and stops. The second and third — the ones that change what game you can play — are invisible to it, because they don’t live in the current workflows. They live in the gap between what your organisation does and what it could do. Efficiency is real. It is also the *smallest* of the three prizes, and the only one the dominant question knows how to look for.

There’s one more thing the work-first frame changes, and for the people in your organisation it may matter most of all. When you start with the work and map capabilities honestly, the human contributions that remain are the specific, named, essential things the work requires and only a human can provide. Chosen, and elevated. Under the replacement narrative, whatever humans still do is implicitly the leftovers — the stuff AI isn’t ready to take *yet*. Same tasks, entirely different story to be living inside.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

The narrative is a management choice

I want to close this chapter by being precise about where the fault sits, because it would be comfortable to blame the technology vendors or the headlines, and it would be wrong. The replacement narrative persists because it's *easy to manage*: it needs no redesign, no capability mapping, no new roles — just a procurement decision and a dashboard of hours saved. The work-first approach persists in being rare for the mirrored reason: it *is* management work. Composing humans and AI on purpose, redirecting freed capacity toward the dormant skills: none of that installs itself.

Which raises the obvious next question: surely somebody sells help for exactly this? What I found is that the market is selling something adjacent — dashboards, registries, control planes, oversight for your growing AI workforce. Watching, in ever higher resolution. What none of it does is manage. That gap between supervising AI and managing it is where the next chapter lives.

Take this with you: Pick a few people on your team and ask each of them: Where are you buried in tasks that don't use your best skills? What do you wish your team could do but can't — no time, no expertise, never had the budget? And where does your way of doing something differ most from your colleagues'? You'll hear the three patterns — dormant, missing, inconsistent — within the first hour, and you'll know more about where AI belongs in your organisation than any use-case brainstorm would have told you. Then notice which of the three your current AI initiatives address. If the answer is "only the first" … the narrative has you.

NOTES

1. Deloitte, *State of AI in the Enterprise*, 9th edition (2026); survey of 3,235 organisations across 24 countries. Consultancy research.
2. McKinsey & Company, "Superagency in the Workplace" (January 2025); survey of 3,613 employees and 238 C-level executives. Consultancy research.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART I · CHAPTER 5

Supervising and managing

Every time one of the big platform announcements crosses my screen these days, I think back to the fire drills. You remember the scene from the start of this guide: the leader asking how many AI solutions are running and what they deliver, and the scramble that follows. In those days, Excel was the closest thing we had to an AI workforce registry.

So believe me when I say I'm genuinely amazed at what the market has built in barely two years. Ten years ago I would not have believed this shelf would exist. Different vendors (Workday, Microsoft, ServiceNow, OpenAI, you name them) take different approaches to governing AI agents and combating agent sprawl, the industry's new name for something very familiar to those of us who have been dealing with a growing number of citizen-developed AI solutions.[1]

Depending on the platform and its integrations, the capabilities now available include:

- **Registration:** a list of what exists, each entry with an identity of its own.
- **Permissions:** what the worker may access, and who may direct it.
- **Metering:** sessions, interactions, users, cost and usage over time.
- **Recording:** what went in, what came out, who did what.
- **Detection:** anomalies and threats, increasingly with AI doing the watching.
- **Lifecycle controls:** activating, deactivating and decommissioning workers.
- **Performance evaluation:** assessing outputs and tracking indicators of quality and business value.

These tools are enablers of everything I'm proposing. The management practice Part II describes needs places to live: worker records need a registry to be written into, role definitions need a system of record, reviews need performance, cost and usage data, and a deliberate goodbye needs a decommissioning path. The market is building the infrastructure for AI workforce management, and some vendors are already extending their offerings into areas traditionally associated with managing people.

What interests me is what happens once that infrastructure is in place.

Someone still needs to determine what a worker is for, what good work looks like, whether it is still worth having, who makes those decisions, and what changes for the humans working with it. These are strategic, operational and cultural questions, and I've watched organisations invest heavily in technology without adequately addressing them.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

You can do the lifecycle, the testing and the security work impeccably and still end up with a worker nobody manages.

Nothing in the method depends on which of these platforms you bought, or on whether you bought any. The practices attach to the worker, not to the platform. Where you have the infrastructure, it makes them easier to run, scale and document. Where you don't, a shared document and a recurring calendar entry will carry you further than you'd think, provided the risks and complexity of the worker allow for it.

If you're evaluating these platforms, evaluate them seriously. They solve real problems, and I'm sincerely glad they exist. Some of the vendors themselves acknowledge that the technology needs to be accompanied by leadership alignment, workflow redesign and organisational change.[2]

What I would encourage you to pay particular attention to is which management practices your chosen platform enables, and who in your organisation is actually going to practise them.

That brings us to the distinction the whole method turns on.

Two different jobs

Supervision watches. It answers: what exists, what is it doing, what does it cost, has something gone wrong? Registries, dashboards, alerts and access controls can finally give us continuous visibility into the questions that used to trigger those fire drills. Supervision is genuinely necessary.

Management directs. It answers: what is this worker for, and what does good look like? Which capabilities, human and AI, does the work need, and how are they composed? How does this worker get better? Is it still earning its place, and who decides what happens when it isn't? Management means setting direction, composing teams, developing capabilities, reviewing performance and making lifecycle calls.

Both jobs need doing, and the trap is confusing one for the other.

Imagine every metric showing green and every identity properly governed, around a worker producing mediocre output nobody defined a standard for, in a role nobody deliberately designed, toward a goal nobody clearly articulated. Even a sophisticated performance dashboard can struggle to reveal a problem when the organisation has never defined what success means in the first place.

This is what I call **zombie AI**.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Systems continue to consume licence fees, compute and maintenance resources while delivering value nobody has verified for months, in roles nobody has reviewed. They are not dead enough to switch off, because switching off is a decision and nobody owns it. You can have a zombie in a beautiful registry: documented, identity-governed, fully monitored, and still serving no meaningful purpose.

Abandonment by budget cycle is what offboarding looks like when nobody practises it deliberately: the bulk and panicked version of a decision this method teaches you to make one worker at a time.

The measurement question

Take the analytics, because that's where the two jobs meet most visibly.

Depending on the platform, you can see how many sessions a worker ran, which of its skills were used, how many people engaged with it, what it cost, what went in and what came out. Increasingly, platforms can also evaluate the quality of outputs and track indicators of business value.

That is highly valuable data, and I would want every manager of an AI worker to have access to it.

But to judge whether the worker is doing good work, someone needs to establish what good would be. To decide whether it still earns its place, someone needs to define the outcome it was brought in for. And to act on either judgement (change the brief, narrow the scope, expand its responsibilities, retire it), someone needs the authority and the mandate to do so.

AI can increasingly help with all of these activities. I would actively encourage using it to prepare performance reviews, identify development opportunities and challenge assumptions about a worker's contribution.

What the organisation cannot delegate is the accountability for those decisions.

Even if a platform has access to your entire corporate strategy, it remains your organisation's responsibility to decide which objectives to pursue, which trade-offs to accept and what constitutes a worthwhile contribution from its workforce.

Better analytics make better management possible. Whether that possibility translates into practice depends on the organisation.

This is also why I find the familiar concept of the "human in the loop" insufficient as a guiding principle for a combined workforce. Where its implementation reduces humans to checking AI-generated outputs, we have placed them in a predominantly supervisory position, stationed inside a process and watching it pass.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

I am arguing for a broader human role: leading the workforce and setting the direction that the loop serves.

That holds even when AI does the watching. Guardian agents can supervise other agents, and analysts expect this to become a significant part of the agentic AI market.[3] But delegating supervision to another AI worker does nothing to change where accountability for the outcome ultimately belongs.

And the humans?

There is another question I pay particular attention to when looking at AI workforce management platforms: what does all this mean for the humans engaging with AI?

I have made it a cornerstone of my work.

Look at the potential asymmetry we're building. Your AI workers are on their way to being meticulously documented: capabilities catalogued, autonomy levels defined, performance traced, costs increasingly visible.

And the humans working alongside them?

The colleague who's asked to collaborate with the newsletter agent: does he know what it's for, when to trust it, when to challenge it, and what his own role in the collaboration is?

The manager listed as "owner" in the registry: did anyone equip her for that role, give her time for it, teach her what reviewing an AI worker even involves?

I keep encountering a version of this imbalance in organisations: the AI side of the partnership gets a system of record and the human side gets a mandatory e-learning.

If we continue down that path, we risk developing better visibility into what our AI can do than into what our people can do. I'll return to what that blind spot costs in the final part of this guide.

The research on human oversight gives us good reasons to take this seriously.

In 2022, Ben Green published an analysis of policies requiring human oversight of government algorithms. His criticism was that these policies frequently assumed humans could effectively oversee algorithmic decisions, despite substantial empirical evidence showing how difficult that can be. Humans can defer to the machine, exhibit exactly the automation bias they're supposed to counteract, and inadvertently lend legitimacy to systems whose underlying accountability problems remain unresolved.[4]

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Subsequent interdisciplinary research identified four conditions for effective human oversight: the power to intervene, access to the knowledge needed to understand the situation, the ability to act on one's judgement, and intentions appropriate to the oversight responsibility.[5]

My reading of those conditions is straightforward: effective oversight has to be a designed job.

There is also a longer-term problem. In a July 2026 article for the World Economic Forum, Isabell Steidel and Benedikt Gieger gave it a particularly useful name: the **oversight paradox**. The expertise humans need to assess AI's work is maintained through practice, so delegating more of that work to AI risks eroding the very competence that oversight depends on.

Aviation has long documented a related problem, automation bias. In flight-simulator experiments, professional pilots often missed events that the automation failed to flag.[6]

Think about what these two problems could mean when combined. We ask people to supervise systems they increasingly rely on, while potentially reducing their opportunities to practise the skills that would allow them to recognise when those systems are wrong.

Read those findings the way I do, and they become a job description for management.

Naming a human in a registry establishes who is supposed to carry responsibility. Making that human able to carry it requires considerably more: equipping them with the right knowledge, giving them time and authority, and deliberately maintaining the expertise their judgement depends on.

I call this **accountability enablement**, and it's so central to the method that it becomes one of the six leadership vital signs in Part III.

The same consideration applies to every colleague working with an AI co-worker daily. A 2024 meta-analysis of human-AI combinations found that, on average, human-AI combinations performed worse than the better of the two working alone, though results varied with the type of task.[7]

For me, that makes the deliberate design of human-AI collaboration one of the most important management tasks. Role clarity, appropriately calibrated trust, and the ability and permission to challenge the machine are conditions I would want leaders to actively create, rather than hope will emerge as people use AI more frequently.

These are deeply human and organisational questions, and no amount of agent governance will answer them on its own.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

From infrastructure to practice

Part I ends on good news. Something has changed, and the change is worker-shaped, though never person-shaped. The worker is a unit you deliberately define, doing work whose division across your combined workforce deserves honest redesign. And the infrastructure for handling much of this now exists: buyable, real, better than anything I could have imagined during the Excel years.

What remains is the work of establishing a management practice around it: setting direction, designing the collaboration, developing both sides of the workforce and creating a culture in which humans understand their roles and are equipped to carry them.

Technology can support that practice considerably. Establishing it, however, is an organisational responsibility, and it can be built on disciplines your organisation already knows.

Part II is that practice. I call it the One-Workforce Method, because that’s the whole claim: humans and AI, one workforce, one management discipline.

And it starts exactly where you’d start with a human: before the hire, with a job description.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Take this with you: Run a fifteen-minute audit of whatever AI oversight you currently have in place: the registry, the dashboards, the access controls, the review board, the mandatory training.

Mark each element W or M, or both where appropriate:

W: Watching. Does it detect, log, alert, restrict or provide visibility?

M: Managing. Is it part of an active practice that sets direction, defines good performance, develops capabilities, reviews value or makes lifecycle decisions?

Pay attention to the difference between having a feature that could support management and having an actual management practice that makes use of it.

Then pick your most important AI worker and name the humans who engage with it, the users and the owner.

Has anyone told them what their role in that collaboration is? Have they been given the knowledge, time and authority to fulfil it?

If the W column is full, the M column is sparse, and those last questions draw blank looks, you have a good indication of where your next investment needs to go.

You may already have much of the technology you need.

NOTES

1. Examples include Workday's Agent System of Record, Microsoft Agent 365 and Entra Agent ID, OpenAI Frontier, and ServiceNow's agentic workforce management, all announced in 2025 and 2026. Forrester's predictions for 2026 also expect the leading HCM platforms to add capabilities for managing digital employees. Vendor announcements and analyst forecast.
2. OpenAI, "Introducing Frontier Alliances" (23 February 2026). Vendor publication. See also The Futurum Group, "OpenAI Frontier: Close the Enterprise AI Opportunity Gap, or Widen It?" (9 February 2026). Analyst commentary.
3. Gartner, "Gartner Predicts that Guardian Agents will Capture 10–15% of the Agentic AI Market by 2030" (11 June 2025). Analyst forecast.
4. Green, B. (2022). "The Flaws of Policies Requiring Human Oversight of Government Algorithms." *Computer Law & Security Review*, 45, 105681. Peer-reviewed research.
5. Sterz, S., et al. (2024). "On the Quest for Effectiveness in Human Oversight: Interdisciplinary Perspectives." *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT)*. Peer-reviewed research.
6. Steidel, I., & Gieger, B. (2026). "The oversight paradox: Why human control over AI may be eroding the very competence it requires." World Economic Forum, 2 July 2026. Expert commentary. On automation bias in aviation, see Mosier, K. L., et al. (2001). "Aircrews and Automation Bias: The Advantages of Teamwork?" *The International Journal of Aviation Psychology*, 11(1), 1–14. Peer-reviewed research.
7. Vaccaro, M., Almaatouq, A., & Malone, T. (2024). "When combinations of humans and AI are useful: A systematic review and meta-analysis." *Nature Human Behaviour*, 8, 2293–2303. Peer-reviewed research.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART II

The System

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART II · CHAPTER 6

Hired like a worker: the job description

In June 2026 I gave a keynote in Vienna about leading when the ground is moving. I talked about trust and leading through transformation, big themes, carefully built. And afterwards, in the hallway conversations and the messages that followed, one idea kept coming back. Of everything I'd said on that stage, the thing people kept referencing was this: *write a job description for your AI.*

It genuinely surprised me. A job description is about the least glamorous instrument management has ever produced. But the more I've sat with that reaction, the more I think the audience understood something precisely: of everything I said that day, this was the piece they could *do*. Straight away, without waiting for a budget or a platform. And, as I'll try to show you in this chapter, that small act quietly answers most of the questions Part I left on your desk.

The four conversations

A job description for an AI is a forcing device for four conversations that, in many organisations, never happen before an AI starts work.

The first is what this AI is hired to do. The vendor pitch answers a different question: what *can* it do? The honest answer is usually "an enormous amount." The JD answers what job we are giving it *here*. A tool might be capable of fifty things. The job description defines the three it is hired to deliver on this team. That distinction, capability versus role, is the unit thinking made operational, and it is the first thing that separates a managed worker from a subscription with ambitions.

The second is what good looks like. Before deployment, define how you will know whether this worker is performing: a few concrete indicators and, this part gets skipped constantly, the *baseline*. What does the "before" state actually look like? First drafts currently take two days; rework on outsourced content runs forty-five percent. Without the before, the after is forever unmeasurable, and the performance review will have nothing to stand on.

And what if there is no before, because the worker does something your organisation has never done? Remember the expansion play: the capabilities that were always parked under "if we ever have budget, we'll hire an expert." Say a team hires an AI worker to synthesise market intelligence nobody ever had the time or expertise to produce. There are no hours

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

saved to count, because the hours were never spent. Then the honest baseline is "this didn’t happen at all," and the success question changes shape. It moves from faster-than-before to delivered-at-a-defined-standard: what does a good market brief contain, which decisions must it be strong enough to support, and who is qualified to spot-check it? That last question deserves real thought, possibly external help at the start, because work nobody in the team has done before is work nobody in the team can eyeball for quality. It is harder to write than a speed target, and more necessary, precisely because the expansion play is where the transformative value lives.

The third is who is responsible. Name a human who steps into the manager role for the AI worker. One name: one person managing it day to day. It is the line most frameworks skip, and the one that decides whether anything else in the document gets acted on.

The fourth is what it means for the humans working with it. An AI co-worker never joins an empty room. Someone’s work feeds it, someone reviews what it produces, and someone’s role has just changed shape. The JD names who will collaborate with this worker, what changes in how their work is currently done, and what they will need to do it well. Of the four conversations, this is the one I have made the heart of my own practice, and the ripple section below is devoted to it.

The fifth conversation

There is a fifth conversation hiding inside the other four: *is there actually a job to be done?* Writing a JD is a sanity check. Sometimes the act of articulating the role reveals that the use case is vague, the expected value thin, the team not ready, and the honest outcome is "not yet." Discovering that at the JD stage costs you an hour. Discovering it after procurement or three months of development costs you a line in the abandonment statistics. That is why the JD belongs *upstream* of the build-or-buy decision. Written first, it becomes the requirements brief that procurement evaluates vendors against, or the specification your own team builds the custom solution from, instead of a rationalisation written after the contract is signed or the build is done. The role decides the technology. Define what you need, then build or buy.

Why writing it down works

I want to be precise about why a written document does what good intentions don’t, because the reason shaped my whole approach to this practice.

One of the first things I learned studying communication at university was that people read the same sequence of events differently. What’s razor-sharp in your head can land blurry, or simply wrong, with the person next to you. An unwritten expectation about an AI’s role lands

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

differently with every single person on the team. The writer assumes the AI drafts and she edits. Her colleague assumes it is there to replace the drafting entirely. The manager assumes someone is checking the outputs. Nobody is. Subjective interpretation is bloody dangerous, and a JD is the instrument that pins the answers down, so everyone works from the same understanding instead of their own private version of it.

What goes in it

The template I use has three layers. The separation matters more than the sections: the role, the technology and governance around it, and the humans.

Layer A is the management view. It is written by the person who will manage the worker, in management language. It is the minimum viable job description, and for a small organisation it may be all you ever need. It covers the worker’s *identity*: a role name, deliberately separate from the underlying tool, because the role persists even when the tool changes, and the one person who manages it. Your "Content Drafter" might run on one model today and a better one next year, and that is why the role, not the technology, is the unit. Its *purpose* names the problem it helps solve or the opportunity it taps into, and the strategic goal it serves, which keeps your AI workers focused on what matters. Its *contribution profile* lists the tasks it carries out and, just as important, the ones it does not. These deserve careful thought, because this is the part that becomes the requirements brief for building or buying. Its *success criteria* start with a baseline, where things stand today or the opportunity you are missing, followed by indicators of what should improve. They also hold something human JDs rarely dare to include: reassessment triggers, the conditions under which this role should be reconsidered. And they set a review rhythm, so "deployed and never evaluated" cannot happen by drift. How often you plan to review the worker is also your estimate of the management effort it will need.

Layer B is the technical and governance view. It typically covers the data, integrations and tools, tool or platform considerations and constraints, budget considerations, and the risk profile and other governance requirements, adjusted to your organisation’s own processes. Support functions help complete it, and organisations with real regulatory exposure will need it. Treat that list as a coverage check. If your organisation already runs data-protection assessments, ethics reviews, procurement processes or a development lifecycle with its own gates, Layer B is not there to duplicate them. Connect it to them instead. A question an existing process already answers gets a reference. A question no process asks yet gets answered here. The point is simply that every one of these questions is answered somewhere, and the JD shows where. The part I enjoy telling sceptical compliance teams: a well-completed Layer B hands them most of what their own assessments would ask for anyway, gathered once, in one artefact, before deployment, when constraints are cheap to accommodate. Done right, the JD is the intake that makes the rest of the pile faster. And if

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

they would rather answer those questions in their own systems and processes, the JD connects to them and brings the answers to the manager, so they find their way into how the worker is actually run.

Layer C is the human view, and if you remember one thing about this template, let it be that it has one, because no vendor's agent-configuration form ever will. Layer C answers the fourth conversation in writing. Which roles collaborate with this worker? What is the collaboration model: AI drafts and human refines, human directs and AI executes, or co-creation, where the human and the AI build the output together in short rounds, each adding to and reshaping what the other produced? What is the decision scope (advisory, augmentative, or autonomous within defined boundaries), and where exactly does work hand off between AI and humans? What changes from how the work is currently done? Which parts of whose role shift, and what do their adjusted roles look like? What skills do the collaborators need to work with this AI effectively: briefing and directing it well, assessing the quality of its output, giving feedback that improves the next round, and the judgment call of when to rely on it and when to override it? And what does it demand of leadership: the capabilities the manager needs to run this worker, such as reviewing performance, spotting drift and AI literacy at the depth the role requires, plus the honest time budget for doing so. Management costs hours. Pretending otherwise is how oversight decays into rubber-stamping. Layer C is written by the manager with the team. That co-writing session is quietly the first act of onboarding.

Scale changes how Layer C gets written. A worker serving a single team has one set of collaborators, and naming them individually is the honest thing to do. The protocol writer serves every team in the organisation, and no single document can hold forty versions of whose role shifts. There the JD names the collaborating roles once, and the ripple conversation happens per team, run by the manager who owns that team's work. One worker, one job description, as many role-change conversations as there are teams.

One principle governs who holds the pen: the manager writes the JD. Decades of HR practice converge on this. The person who will manage the work defines the role, while support functions advise on feasibility and constraints. Role definitions written by support functions alone come out technically correct and organisationally disconnected. It is as true for AI briefs written solely by IT as it ever was for human JDs written solely by HR. The one exception is deliberate: Layer C is co-authored with the people it is about.

In my own practice, managers draft the JD with an AI assistant at their side, and its role is deliberately narrow. It plays the role HR holds for human JDs. It doesn't suggest what the role should be, propose success criteria or fill in a single answer. Every piece of content comes from the manager. What the assistant does is what a good HR partner has always done in the JD process: hold the structure, make sure every section is covered and, when

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

the manager gets stuck, help them find the right question to ask themselves, never the answer. At the end, it turns the conversation into a clean, structured document. If you are smiling at the irony of an AI helping write the job description for an AI, so am I. But notice the design: even this assistant has a defined role with clear boundaries. The thinking about what the role is for, the decision to put one name on the accountability line, the conversations with the people whose work is about to change: none of that ever leaves the manager's desk.

The move that makes this chapter matter: the ripple

The part I consider most important in this chapter: the JD is the first step of the whole method, and this is why.

A job description for an AI is really a document about your people.

That is why the humans get a layer of their own. Inside Layer C sits the question I consider the single most valuable in the entire template: *what will people do differently once this AI is on the team?*

Answer it honestly and something remarkable happens: you discover you are rewriting several job descriptions at once. If the AI takes over first drafts, the writers' roles have shifted from drafting to editing and direction. That is a different job, with different skills that deserve naming and developing. If the AI summarises every meeting, the junior colleague who used to do it needs a new path to visibility. In my experience this ripple is where the actual transformation happens: the honest mapping of who does what, executed one role at a time. Organisations that deploy AI without it are doing precisely what the returns data in chapter 4 showed, adding capability while leaving the work redesign undone, then wondering where the value went.

The ripple is also where trust is built or broken, and this is what the Vienna audience heard in the idea. When an AI joins a team with no defined role, your people fill the silence with the worst interpretation, usually some version of "this is here to replace me." You cannot trust what you cannot define, and neither can the people working next to it. A written JD replaces the silence: what the AI is for, where its work ends, who oversees it, and, named explicitly in writing, what the humans around it contribute and what new opportunities their shifted roles open up. Skip it, and the AI quietly becomes the hero of every success story while the humans start to look like overhead. That is never the message you want to send to the people you are asking to trust you. If you care about your people, write a JD for your AI.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

"Isn’t this bureaucracy?"

I hear the objection, and I have learned to welcome it, because the answer is where this chapter’s practice earns its keep. Yes, it is a document, and the corporate immune system rightly attacks documents that exist for their own sake. So hold the JD to the same standard: every section exists to force a conversation that would otherwise not happen. If a section doesn’t force a useful conversation, it does not belong in the template. One page for Layer A, written once and updated at reviews. The paperwork itself can go to the assistant described above. Against that sits the fully loaded cost of one zombie AI, a worker nobody manages but everybody still pays for: licence fees, integration upkeep, the slow corrosion of a team working next to an unexplained machine, and the eventual write-off when the budget cycle finally notices. Writing a JD for your AI is a cheap thing to do, and an expensive one to skip.

You wouldn’t keep someone on your payroll with no clear role in the business. An AI without a defined purpose is exactly that: burning money, unsettling your team, adding nothing but plausible-looking output. The JD is how you make sure you never hire one.

And once it is written, it becomes something more: the first entry in the worker’s file. Every subsequent step of the method (onboarding, reviews, the eventual goodbye) will add pages to that file. Which raises the next question. The JD says what this worker is hired to do. How does it actually *join* the team? Deployment is a technical event. Onboarding is an organisational one, and that is next.

Take this with you: Pick one AI already doing real work in your organisation (the newsletter agent, the meeting summariser, whichever you know best) and write its job description this week. The essentials fit on one page, six questions: What is it hired to do, the three things it delivers here? What does good look like, and what is the baseline? Who is the one named accountable human? Who works with it, and what changes from how the work is done today, named specifically? What do those humans need to learn to work with it well? And when is the first review? Write the last three questions with the team. Half of what you write will trigger a conversation you realise is overdue. That conversation is the method working.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART II · CHAPTER 7

Onboarding: the AI worker file

A new AI system arrives in most organisations like this: The licences go live. There's an announcement in the next team meeting, maybe a link to a training video. IT confirms the integration works. Somebody posts a cheerful message with a rocket emoji. Done, we're live. If you're one of the good ones, there's more: live training instead of the video link, office hours in the first weeks for the questions that only surface once people actually try it, maybe a few power users seeded across the teams to help their colleagues along. Noticeably better. And still, look at what all of it is: teaching people to operate a tool.

Now watch how a new colleague arrives, even in organisations that do it badly. Someone meets her at the door. She gets introductions: this is who you'll work with, this is who to ask when you're stuck. Someone explains the unwritten rules, the house shorthand, and where the real information lives as opposed to where the org chart says it lives. Her manager sits down with her: here's what the first weeks look like, here's what good looks like, here's when we'll talk about how it's going. Nobody would call activating her badge "onboarding."

Onboarding runs in two directions

The mistake in the rocket-emoji rollout is thinking there's one thing to onboard. There are two.

You onboard the worker into the organisation. A new human hire arrives with her capability. What she doesn't have is your context: the examples of what good looks like here, the client's history, the unwritten rules, the shorthand. Good onboarding transfers exactly that. An AI worker needs the same transfer, and needs it more urgently, because it will never pick your conventions up by osmosis at the coffee machine. This is the briefing work. The job description identified the context this role needs; onboarding is where that context actually reaches the worker. The test is the one you'd apply to a capable new hire: what would she need to read, see and be told before her first piece of work is any use to anyone? Some of what the worker needs to read sits behind a login, which is where onboarding turns into a permissions question. Skip it, and you get the thing I warn every client about: a generic capability bolted onto your people's workflows, producing plausible output that belongs to no one's actual work.

If your worker was built rather than bought, this transfer may already have happened during development. In a well-run project, the team doesn't just construct the retrieval system. They curate what goes into it together with the people who know the work, the subject-

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

matter experts. That is precisely the context transfer this section describes, done early and built into the worker from the start. If that's your project, onboarding doesn't repeat the work. It confirms it's covered, the same principle the job description applied to governance: answered somewhere, checked here. Where it hasn't happened, because the worker was bought and configured out of the box, or because the build treated context as a data-plumbing task rather than a conversation with the people who know the work, then this is the time. (What that context consists of, and who keeps it current, gets its own treatment later in the guide.)

The job description already says what this worker needs to reach. Onboarding is where that becomes real: granted deliberately, scoped to the role, and revoked when the worker changes or goes, the same reflex a departing human employee triggers without anyone having to think about it. Access beyond the role is exposure with no upside.

And you onboard the organisation to the worker. This is the direction almost everyone skips, and it's where the ripple becomes operational. The humans who will work with this AI need to know the rules of collaboration: what the AI does, where its work ends, who checks what, where the handoffs sit, and what to do when an output looks wrong. The escalation path needs a name and a channel, agreed before the first incident, when there's still time to think. And the role shifts the JD identified (the writer who now edits, the junior who no longer takes the minutes) need to be talked about as growth conversations, not discovered as quiet losses. Whether human-AI collaboration works is decided here, in these unglamorous agreements.

The introduction

Between the licences going live and the first review sits a meeting that costs half an hour and that hardly anyone holds. The manager introduces the new AI co-worker to the people who will work with it.

Most organisations will say they already do this. What they usually do is a demo. Someone presents the new system, shows what it can do, walks through a few prompts, and takes questions at the end if time allows. Everything in the room points one way: from the front. And since whoever is presenting usually built or bought the system, the tone is proud. Capabilities lead, limitations arrive as footnotes, and the team leaves with expectations no worker could meet.

Overselling at the introduction is not a small sin. It sets the baseline every later output will be measured against and fall short of, and it teaches people that the honest conversation about this AI co-worker happens somewhere other than the room they were invited to.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Be precise about what that meeting is ... because it sits one word away from something ridiculous. It is not a welcome party for software. It is a manager telling a team that their work is about to change, and saying what is changing while there is still time to ask about it. The AI is the subject of the meeting. The people in the room are the point of it.

In my own practice it runs roughly like this. What this AI co-worker is here to do, and which work it takes off whose desk. What it does badly, said out loud by the manager before anyone discovers it the hard way. Where people will need patience, and for roughly how long. Who to go to when an output looks wrong. Then questions, for as long as there are questions. The questions earn the half hour on their own, because they surface every assumption the job description never thought to make explicit.

Then comes the move that changes the temperature: invite feedback from the first day. Small corrections are cheap while the collaboration is still forming, and a team that was asked at the start tends to speak up later. That's how you hear about the quiet workaround before it hardens into how things are done here.

None of this survives being imported wholesale into a culture that doesn't work that way. An organisation that never introduces human joiners with any ceremony will find it odd coming from an AI. Run it in whatever form your people already recognise for a change to their work: a slot in a meeting that exists, or a short round in the team channel. The form is negotiable. What matters is that the human manager of the AI worker does it, in front of the team, and doesn't just send a link.

The AI worker file

Human onboarding has an artefact at its centre: the personnel file. It doesn't start empty. The role existed before anyone was hired, and on the first day it is joined by everything the arrival itself produces: the signed contract, the start date, the manager's name, the reporting line. From then on it never stops growing.

Your AI worker's file works the same way, and you have already written its first entry. The job description gets its second half the day the worker actually joins, when the build reaches a version people can work with or when the licence is signed. What joins it is what only now exists, in three parts. The job description is the first. The second captures the technical and governance records: documentation, versions, integrations and where the system metrics live, along with the risk assessment, approvals, data protection and other compliance documentation. Both are referenced where they already sit, or kept in the file if they have no other home. The review record is the third. It opens with the rules of collaboration agreed with the humans who will work with the worker, and from go-live onward it holds every check-in, role change and decision, through to the goodbye.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

That living document is the **AI worker file**, and it is the backbone of the management practice this guide describes. Everything Part II asks of you is either written into it or read out of it.

Before you object that this sounds like a new system to maintain: it isn't one. A Confluence page or a folder with a template will do, and so will an entry in whichever registry the market sold you. The file flexes to fit how your organisation already documents things. The discipline is in the thinking. The file compounds by design: every review adds a record, every role change adds a version, and when the day comes to say goodbye, the file is what makes the goodbye clean. Multiply it across workers and you get the portfolio view a leadership team needs when the fire-drill question arrives, answered from information already at your fingertips.

The first weeks

One more habit to steal from good human onboarding: the early check-in. No sensible manager hires someone, sets objectives, and walks away until the annual review. The first weeks are when expectations meet reality. You find out the brief was ambiguous, the handoffs sit in the wrong place, the team has questions nobody anticipated. The file therefore carries the date of the first check-in after go-live, a few weeks in, set by the review rhythm in the JD.

Early reviews are short and frequent, and mostly about the collaboration. Is the output landing? Are the humans using it as intended, and is anything drifting? How you run them should fit how your organisation already works: a slot in a meeting that exists, a quick round in the team channel, or an automated check-in agent that asks every collaborator a few short questions and consolidates the feedback for the manager. (Yes, an AI helping to review an AI. The division of labour is the same as with the JD drafting assistant: it gathers and structures, the humans supply the answers, and the manager supplies the judgment.)

No more than five minutes of anyone's time, in whatever form. Those five minutes matter to whether the worker delivers, because they catch drift and quiet non-adoption while there's still time to fix them. As the worker stabilises, the cadence relaxes. That's the manager's call, based on risk and maturity rather than a fixed corporate calendar.

Build problem or management problem

One clarification belongs here, because leaving it implicit costs me the engineers in the room every time. Saying that deployment is not onboarding is not saying that deployment doesn't matter. The build discipline (the development lifecycle, testing, evaluation runs, technical documentation, security review, and monitoring that tells you the thing is still

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

standing) is essential, and this method leans on it. The job description's governance layer connects to those processes, and the AI worker file points straight at their artefacts. A badly built worker cannot be managed into performing. I have never once seen it done.

What I am claiming is narrower, and organisations keep discovering it the expensive way: **engineering decides whether it works. Management decides whether it performs.** Two conditions, both necessary, neither able to stand in for the other, and only one of them has a mature profession behind it. Software engineering has spent decades building standards, roles, review gates and tooling for the first condition. For the second, most organisations have roughly this: a shelf of instruments (if at all) and no practice.

Put it in human terms and it stops being controversial. A new colleague can be recruited flawlessly, contracted correctly, equipped on day one and cleared by every compliance check, and still deliver nothing for six months, because nobody told her what good looks like here or sat down with her until the annual review. When that happens, nobody blames the payroll system. We know precisely which of the two conditions failed. With AI workers that clarity disappears: the system underdelivers, and the reflex is to send it back to the technical team for a better model, more retrieval or another integration. That reliably produces a more sophisticated version of the same disappointment.

So when an AI worker disappoints, the first question is a triage question, and it takes about a minute. Is this a build problem or a management problem? Wrong answers, retrieval that misses, a regression after an update, or a system that falls over under load are build problems, and they belong with the people who own the build. Technically correct output that nobody uses, output that serves no outcome anyone can state, a worker nobody has looked at since go-live, or a capability that quietly duplicates what another team already runs are management problems. No amount of engineering will fix those, because nothing is broken. Both kinds are real and both are common. They have different owners and different remedies, and the most expensive habit I meet in the field is routing the second kind to the first kind's team.

That leaves the question the whole method has been building toward: what does reviewing an AI worker's performance actually involve, and why do the dashboards tell you so little about it?

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Take this with you: Complete the file. The job description is its first entry. Add the technical and governance records from the build or the purchase, the rules of collaboration, and the date of the first check-in after go-live, and put the file wherever your team already keeps documents. Then decide how you will introduce this worker to the team, in a format that invites open discussion, and think through how you will manage expectations: what you'll admit it can't do, and where you'll ask people for patience.

PS: it's never too late. For an AI that has been running for months, the file and the open discussion work exactly the same, and both tend to surface what nobody knew.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART II · CHAPTER 8

Performance review

Your AI worker has been running for three months. Is it doing a good job?

In most organisations nobody knows, and not because the data is missing. The platform will tell you how often it ran, who used it, what it cost. None of that says whether the work was any good. That judgement needs something to measure against — and you already wrote it down: what this worker was hired to deliver, what good looks like here, and when the first review would happen. The AI worker file has been waiting for this since go-live.

Monitoring tells you the system runs. Reviewing tells you the worker delivers.

The distinction holds here too: the shelf watches, management directs. Performance is where the two get confused most expensively. Monitoring — uptime, latency, usage, drift metrics, cost — is supervision applied to performance, and you need it. But monitoring observes the AI *alone*. The unit that produces value is the combined system: the AI worker, its context, and the humans around it.

What the review hands on

One thing a review of an AI worker reliably produces is a finding about the humans. Outputs going unused, briefings arriving vague, an AI worker whose quality nobody on the team can properly judge — none of those is an AI problem. Each is a development need: a skill to build, an input to improve, a role shift that was never really talked through. The review's job is to surface it and write it down. What happens with it belongs to whoever owns that person's development, and to the chapter on leading a human-AI team.

Every review ends with a decision

A review also has to be honest about cost, which means counting what keeps accruing after procurement: subscription and token spend, yes, but also context maintenance, integration upkeep, and — the one everyone forgets — the human oversight time the AI worker consumes. A leader who measures AI cost against the original purchase price will systematically undercount. The comparison that matters at review time is the one the job description wrote down: what would this outcome cost without the AI worker, and is the answer still favourable?

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

The discipline that turns reviewing from observation into management, and the single sentence I'd have you keep from this chapter: **every review of an AI worker ends with a lifecycle decision.** Four options, the same four you'd recognise from managing people: *Grow* — the AI worker earns an expanded mandate, more investment, new capabilities. *Maintain* — continue as is, revisit next cycle. *Re-skill* — the role has shifted; back to a revised job description. *Retire* — the worker no longer earns its place; begin the structured goodbye. The decision goes in the AI worker file, with the review record.

Practicalities, briefly. Cadence is the manager's call, set by risk, maturity, and criticality — a customer-facing AI worker in its first month might need weekly attention; a stable internal AI worker, anywhere from quarterly to yearly. The format can be light: a structured check-in between manager and collaborators, fifteen minutes, four questions — performance against targets, collaboration quality, cost versus value, risks and drift — captured in the file. It can even be partly automated: a short conversational check-in agent that gathers collaborator input and updates the file, the useful face of AI supervising AI — leverage for the human reviewer, never a replacement for the decision, which stays with the named manager. And the reviewer needs to know what AI failure looks like, because it rarely looks like failure: confident wrong answers, gradual drift as context ages, plausible output that quietly stopped matching reality. A manager who has never seen an AI break will read fluency as health — one of the leadership capabilities Part III takes up properly.

One decision from the four deserves its own chapter, because organisations find it hardest, skip it most, and pay most for skipping it. The goodbye.

Take this with you: Run one fifteen-minute review this week, on the AI worker whose file you opened. Four questions with the collaborators: Is it delivering against the targets in its file — measured, not felt? Is the collaboration working, and where does it chafe? What does it actually cost, oversight time included, and is that still worth it? Anything drifting — context ageing, confident errors, workarounds creeping in? Then close with the sentence that makes it a review: "Decision: grow, maintain, re-skill, or retire — because …" Write both the answers and the decision in the file. If the decision was hard to say out loud, you have just learned the most important thing in the review.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART II · CHAPTER 9

The structured goodbye

When someone's work is no longer needed, you part ways. You would never let a colleague who has left keep their login and their access to your data. That would be ridiculous.

Human workforce management has built a whole discipline around what it calls offboarding, with handovers and access revoked on the last day.

With AI we are far less disciplined. Partly because AI workers aren't physically there. You don't see them hanging around the office doing nothing, so there's no urge to act. And partly because switching off an AI worker can feel like admitting failure. Someone championed it, budget was spent, maybe careers were attached to it. So the rational move for everyone individually is to let it fade quietly rather than end it visibly. That is exactly the mechanism that fills portfolios with AI zombies: those invisible, unused agents and automations that add zero value but still produce cost and risk. An AI worker that adds cost and risk without adding value is a liability.

So in addition to the first step, which is that you have to take a deliberate decision to retire an AI worker (and this happens through the review process), you should also consider a few things when you switch one off.

Who picks up the work? The AI worker delivered a distinct outcome. That outcome either transfers to a human, moves to another AI worker, gets absorbed into a redesigned process, or is consciously discontinued. Consciously is the word: the job description tells you exactly who and what relied on this AI worker, so nothing falls through the cracks and no downstream team discovers the gap by accident weeks later. And if the work transfers to a person, their role changes with it, and that change belongs in their work and growth plan.

What knowledge must survive? That might be outputs worth archiving or data with retention obligations. The context is usually the part that matters most. The curated briefing this AI worker ran on is organisational knowledge that outlives the worker: the standards it was briefed on and the ways of working written into its system instructions. These are probably the most valuable things the worker leaves behind, and the head start for whatever replaces it. Check before anything is deleted: a reference file or a prompt this worker used may well serve others too.

What gets switched off, completely? Licences need to be cancelled, and integrations and access need to be cut. A retired-but-connected AI worker is the worst of both worlds: no value, but a full attack surface. Every forgotten system is an unlocked door that nobody is watching anymore.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

What did we learn? Before the file is archived, the manager writes the learning record: what this AI worker was hired for, what it actually delivered, what worked and what didn't, and what the next job description in this space should know. Offboarding is the organisation's learning moment. Across a portfolio, these records become the organisation's AI learning curve made visible: the patterns behind failures and the practices behind successes.

Retirement is a success

A clean retirement is evidence that the management system works. It means someone made the call at review and carried it out properly. The worker was hired against a need, served it, and was retired when the need moved on. That is a lifecycle completed, and the fact that it happened by decision rather than by budget-cycle cull is precisely the difference between managing a workforce and warehousing software. I encourage clients to notice first retirements the way they would notice first hires: this is the discipline, visibly alive.

That completes the lifecycle: hired against a written role, onboarded with the organisation prepared around it, reviewed with every review ending in a decision, and retired with intention.

Which leaves the question Part III answers: who runs all of this? People, in roles that are part familiar management muscle and part new. One of them might be you.

Take this with you: Find your most likely zombie: the AI worker nobody has looked at in six months. Give it the review it never had. If the answer is retire, name who picks up the work, check which context it shares before anything is deleted, cut the access, and write the one-paragraph learning record. Do it once and the second one gets easier. Zombie-hunting becomes a habit, and a healthy one.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART III

The Roles

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART III · CHAPTER 10

The manager of an AI worker

The thought at the centre of the One-Workforce Method came to me in the middle of a fire drill. I was collecting the answers to a leadership question — how many AI solutions are running, and what value are they delivering — and I was going through the owner column of the Excel file I had curated. It had everything from business unit leads and product managers to delivery leads. Every row had a name, and the names meant different things. Depending on which part of the business and which use case you looked at, being the owner of an AI solution came with a different description and different expectations. And then there was a name of an owner who had already moved to a different part of the business, who might be able to answer questions about what the solution was doing in the past, but likely had no insight into its current value to the business.

That was the first time I wondered why it had to be that way: a responsibility written down on a list, which some people take very seriously and others may not know they hold.

And I wondered whether it wouldn't be nice to have something like an org chart. Where it is clear who is managing the thing day to day and when people move on, their solutions get mapped to a new manager.

And being a “manager” is a role everyone understands. Nobody has to explain what a people manager is expected to do. That knowledge still holds when the worker is an AI. Much of it transfers as it is, some of it changes. Same but different.

What a good manager does

Defining the role. Every good manager tells a new person what the job is, what good looks like and when they will talk about how it is going. With an AI worker, that conversation becomes the job description, and it matters even more, because an AI worker rarely asks, and it never pushes back. It does whatever it is given. So the role has to be clear before the work starts and kept current afterwards. The manager can turn down new tasks that don’t fit the job description, or change the job description deliberately. What she can’t do is let the role drift because someone found one more thing it could do.

Organising the work. A good manager knows who does what in the team and makes sure the handovers work. With an AI worker, the manager decides what the AI worker does and where its work meets the people around it, from the input it needs to the people who work with its results. She talks to those people regularly about what works and where they have to correct it, because they see things no dashboard shows.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Reviewing and growing. A good manager looks at how the work is going and helps people get better at it. With an AI worker, the review asks two questions on a regular rhythm: is it being used, and does it deliver what we have it for? When the answer is no, the first thing to work out is why. If something is technically broken, that goes to whoever built it — most often the manager herself or a colleague, sometimes with help from a colleague, IT or a vendor. If it works and still doesn't deliver, the problem stays with the manager. Growth looks different too. An AI worker gets better when its briefing gets better: the right knowledge, kept current together with the people who own that knowledge. Sometimes growth means a wider mandate, decided on purpose and written into the job description.

Acting. A good manager acts when something isn't working. With an AI worker, the options are clear: keep it running or sharpen it, pause it when needed, or end it on purpose. She can have its access checked and restricted, and she can demand changes from whoever built it. When responsibilities are unclear, or a problem goes beyond her role, she escalates. All of this needs the authority to decide, which is why that authority belongs in the role itself.

Influence without authority

The manager of an AI worker also has to start thinking about human roles. When the AI worker takes on new tasks, someone's work changes: a step they used to do disappears, a check they never had to do appears. Those changes rarely announce themselves. The manager of the AI worker is often the first to see them coming, and the one who has to bring them early to the people affected and to their managers.

That takes influence more than authority. The manager of an AI worker is not necessarily the manager of the humans who work with it; their roles, goals and development stay with their own people manager. The next chapter is about that manager.

What it does for you

The first AI worker I properly onboarded was a writing assistant, a co-creator for my own posts and papers. I sat down to brief it: how I write, what good looks like. I expected it to take an hour. It took days. It turned out I had rules, about rhythm and about the words I never use, that I had followed for years and never written down. The AI needed them written, so I wrote them. And the file didn't only make the assistant better. Seeing my own habits stated as rules, I noticed which ones were craft and which were just habit. I write more deliberately now than before I taught a machine to write like me.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

That is the first thing managing an AI worker does for you: it makes your expertise visible, to the AI and to yourself. You can't tell an AI worker to "use good judgment", so you have to work out what your judgment actually consists of. And writing a rule down means rehearsing it, which keeps your skills sharp instead of letting them fade as the AI takes over the work.

The second thing is less obvious. Experts often keep what they know to themselves, because it is what makes them valuable. With an AI worker, that logic turns around. Its success is your success as its manager, so every piece of expertise you put into it improves the very thing you are measured on. That holds as long as your own role stays clear, which is why the growth conversation in the next chapter matters.

Who takes this on

Experienced leaders already know how to do most of this. They have written job descriptions and run reviews. They have made the hard calls. What changes is the worker, and with it the details of each duty. The skills transfer.

Some worries fall away: an AI worker has no working hours to respect and no mental health to look after. The care that would have gone there is needed even more by the people working with it, whose work and sense of their own role can shift with every new task the AI worker takes on.

For many others, managing an AI worker will be the first time they manage anything at all. A lot of people will have an AI worker as their first report. That makes this role a first real opportunity to build leadership skills: saying what good looks like and giving feedback that actually changes something.

Take this with you: Pick one AI worker you are responsible for, or one you rely on, and go through the four duties in this chapter: defining the role, organising the work, reviewing and growing, acting. Mark every one nobody is doing today. That list is the job you didn't know you had, or the conversation you need to have with whoever does.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART III · CHAPTER 11

The leader of a human-AI team

A leader once shared his frustration with me. AI was coming into his team from all directions, like new hires who had sneaked in without going through the hiring process every human candidate faces. Some of them changed how the work got done and who did it. Together they changed his team. Getting back into the driver's seat and steering this new combined workforce was, for him, the biggest leadership challenge he had faced in a long time.

A shared north star

Getting back into the driver's seat starts with knowing where you want to go. Why does this team need AI at all, and how does it support what the team is there to achieve? Which values shape how work gets done here, and what do they mean once some of that work is done by AI?

If AI just sneaks in, it won't work smoothly, least of all where it runs against how the team believes work should be done. A team needs a shared north star that helps everyone understand when bringing in an AI worker is a good idea and when it isn't.

Seeing the whole team

To steer a combined workforce, a leader needs to see it whole: the people on the team and the AI workers alongside them, and what each of them can do. That picture is the basis for every decision about who, or what, takes on a piece of work.

For the human side, most organisations already have the tools, from HR systems to capacity plans. For the AI side, the picture comes from the AI worker files. Put together, they become a portfolio view of the team's AI workforce: what each AI worker is there for and how it is performing. That view lets a leader steer the AI side as deliberately as the human side. If you don't have a system for this yet, an Excel sheet or SharePoint list that pulls the latest information from the AI worker files will do. You can even compile it by hand to begin with. Once you have more than five AI workers, it's worth the effort to automate.

Both pictures hold part of the story. The AI worker files show where people are in the loop and what changes for them. The HR systems hold people's roles and development plans. What neither does on its own is connect the two: a change recorded in an AI worker file only

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

reaches someone's goals if somebody carries it across. And some things neither captures well, such as how much of a person's skill goes into making an AI worker's output usable. The leader's job is to bring the two pictures together and act on what falls between them.

Safe to say

Seeing the whole team only works if people tell you what is really going on, and with AI there is a lot they have reason to keep to themselves. Some hide that they use AI at all, because people who use it are often judged as less capable for the same work.[1] Others hide that they don't understand what an AI worker did, or that they quietly overrode it. And underneath sits the question few people ask out loud in a meeting: what does this mean for my job?

One kind of hiding looks like helpfulness. When an AI worker gets something wrong and a colleague quietly fixes it, the problem disappears from view and stays in the system. Thank the colleague, then ask what they fixed and how often.

How the leader talks about AI sets the tone for everyone else. People take their cues less from their own experience with an AI worker than from how the people around them treat it.[2] One senior voice full of scepticism, or one manager trusting it blindly, travels through a team faster than any guideline.

That is why the team needs room to talk about it. Give feedback on working with AI its own room: where AI has helped, and where people had to step in. How you do that depends on your culture and ways of working. It can be a standing slot in a meeting that already exists, or something you set up deliberately. What matters is that the room is there, and that saying the uncomfortable thing in it costs nothing.

Growth and development

A person's role rarely changes because of one AI worker. Most people work with several, some of which they manage themselves and some managed by colleagues. Each of them shifts a little of what the person needs to be good at, and their managers are often the first to notice the shifts. The leader's job is to bring all of it together in one growth conversation with each person, and into a plan for how they can learn what they need.

The conversation fits into the one-to-ones most teams already have:

Which humans and which AI workers do you work with to get your work done? Which of the AI workers are crucial, which are a pain, and which would you wish to have? How is your work changing because of them?

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

The answers show you which growth opportunities and concerns you will have to address. What belongs to the person's development goes into their goals or development plan. They also help steer the workforce: what people say about the AI workers they work with belongs in the portfolio view as much as any performance figure.

Then there is the time that gets freed up. Decide where it goes, into new projects or new learning. Left alone, it is quietly absorbed, and nobody can say afterwards what the AI worker made possible.

And sometimes an AI worker takes over most of a role, and no growth path fits. Leaders have handled roles that were no longer needed before, and the same applies here. Say it early. Where the changes substantially alter what a colleague was hired to do, follow the HR procedures your company has set up for those situations. The procedure stays the same. Only the reason is new.

A check-up for the leader

Leading a combined workforce means keeping several things healthy at once, and it helps to take stock now and then. A check-up measures vital signs one by one and then reads them together, because the pattern says more than any single value. For each one, ask which line describes your team. Count it as healthy only if you could show an example.

Team composition. *Healthy:* for any AI worker on your team, you can say why it is there, and which human roles shifted around it. *Warning sign:* AI added without redesigning the team around it. The humans quietly rearrange themselves around the newcomer, and a year later nobody can say who does what.

Context discipline. *Healthy:* you can name the context your AI workers run on, and who owns each piece. *Warning sign:* context treated as documentation nobody maintains. Without an owner it drifts, and drifting context is the quietest way an AI worker stops earning its place.

Strategic clarity. *Healthy:* you can state your team's direction in two sentences that a person and a system without judgment could both act on. *Warning sign:* leadership describes the strategy one way and people on the ground describe it another, and AI workers get brought on board without goals anyone has written down.

Performance. *Healthy:* you can show both human and AI contributions to your team's results, and you know what your AI workers cost to run in full, including the time people spend checking and correcting their work. *Warning sign:* separate pictures. The AI's dashboard looks green, people's reviews happen somewhere else, and nobody looks at what the team achieved together.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Psychological safety. *Healthy:* in the last month, someone told the team they didn't understand what an AI worker did, or that something wasn't working, and it cost them nothing. *Warning sign:* you are the last to hear when an AI worker fails.

Accountability enablement. *Healthy:* every AI worker has a manager, and when you ask them whether they feel equipped to judge its work and act when something goes wrong, they say yes. *Warning sign:* accountability as paperwork. A name in a field, and everyone knows the named person could not have prevented real damage.

Then read them together. If clarity shows a warning sign, work on it before composition: without knowing what the team must deliver, adding AI turns into collecting capabilities because they are impressive. If psychological safety shows a warning sign, distrust your performance reading, because hidden failures never reach it.

Ask one person on your team to take the same readings of you. Where your answers differ, you have the most honest feedback you will get this year. End with one prescription: the warning sign that costs your team the most right now, treated deliberately for a quarter, with a follow-up date.

Take this with you: Take the three questions into your next round of one-to-ones. Then put what you heard about the AI workers next to your portfolio view. Where the two pictures don't match, you have found your next leadership decision.

NOTES

1. Jessica A. Reif, Richard P. Larrick & Jack B. Soll, “Evidence of a social evaluation penalty for using AI,” *PNAS* 122(19), e2426766122 (2025). Peer-reviewed; in four experiments with over 4,400 participants a social penalty for AI use was revealed: Individuals who use AI tools face negative judgments about their competence and motivation from others.
2. Türkü Erengin, Roman Briker & Simon B. de Jong, "You, Me, and the AI," *Journal of Organizational Behavior* (2025). Peer-reviewed; a human teammate's trust in an AI teammate predicts the employee's trust in it, and this effect vanishes when the human teammate is seen as less trustworthy.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART III · CHAPTER 12

The enabler

Years ago, for about six weeks, I had a data scientist on call.

He worked at a consultancy, and for those weeks, his job included me. I was working on my own use case on a data science platform, I had no technical foundation to speak of, and I could call him. When something broke, or when I had several ways forward and no basis for choosing between them, I called, and he had time.

It remains one of the best learning experiences of my working life, and for a long time I thought it was because of his expertise. But I had worked alongside excellent people for years, before and after, and I have never wanted for colleagues who knew more than I did and were generous with it. Not one of them ever made me feel like a burden.

I still rationed them.

I kept a private ledger without ever deciding to. This person helped me last week, so this week I will work it out myself. That one has a launch coming, so I will not bring her something that takes twenty minutes to explain. I asked the questions that could be answered quickly and I quietly dropped the ones that could not, which, as it turns out, were the questions worth asking. *Talk me through the three ways I could approach this, and why you would pick one.* That question costs a colleague forty-five minutes and costs you a favour you may need later for something more urgent. So you skip it, you pick a path on instinct, and you learn considerably less than you would have.

The six weeks were different for one reason: asking was free. Nobody was doing me a kindness. It was his job, and knowing that changed which questions I brought.

What should concern anyone running an organisation is that none of the questions I swallowed over those other years left a trace. There was no ticket and no line in any report. Everyone involved would have described the situation as working well, and by the measures available, it was. What nobody saw was how I spent my own time: almost two weeks figuring out something that someone could have answered in forty-five minutes. And nobody saw the work I had in mind and never brought to life, because I couldn't get past my own limitations.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Who enables the manager?

When I had my six weeks, I was building machine-learning models, long before generative AI made it possible to ask a machine almost anything. Looking back, many of the questions I once took to a colleague I would ask an AI today. But an AI doesn't know how your company works or what your business needs, so the answer you get may not be the one that helps. And sometimes you don't even know how to put the question in the first place.

I see the same thing now, as more and more people are empowered to build their own AI workers with no development experience at all. They don't know what to watch out for, or how to test and refine what they built until the output is good enough. The obvious response is training, and training helps. But capability rarely transfers in the room. It transfers afterwards, in the first real attempt, at the moment the thing doesn't behave the way it did in the exercise. That moment is where people either build something or quietly give up, and what decides it is whether there is somebody to ask.

The check-up put it as a reading: if the people you named as managers don't feel equipped, you have an enablement gap you need to close. Which raises the question this chapter exists to answer: who enables the manager?

Ideally the person managing an AI worker can build and adjust it themselves. Realistically, hardly anyone starts there. They grow into it at very different speeds, and in the meantime they are accountable for something they cannot yet fully operate. That produces anxiety more than motivation, and a quiet kind of paralysis where nothing visibly fails because nothing much is attempted.

One function, different forms

The function, stated so it holds in any setup: **somebody is dedicated to helping people who are building and managing AI workers, in a way that costs the asker nothing to use.**

The form varies with the size of the organisation and its setup. In the organisations I work with, the role exists, sometimes as one person and sometimes as a small group. It can be a full-time role or a protected share of someone's week. And part of it can be carried by an AI worker built for exactly this purpose: an AI enabler that answers people's questions about building and managing AI workers in your particular setup. It gets its own section below.

Whoever fills the human part, look for someone who learns continuously and in public, and who is comfortable saying "I don't know, let's find out". What your AI platform can do changes within months, so today's expert is out of date by next year, and an enabler who protects an authority position slows down exactly the experimentation the role exists to accelerate.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

And fill it from inside. Even where an outside person covers the role at first, the aim has to be handing it to someone internal. As long as people depend on someone external, there is always a barrier between them and the help they need, whether it is cost or availability. The private ledger comes back, with an invoice attached.

I first saw that this is a role of its own in a champions network I led. The network worked, and it was full of capable, willing people. We still ended up forming a smaller group inside it whose focus was helping others when they dared to build. Helping somebody through their own build turned out to be a different job from championing. If you have a champions network and wonder whether you also need this: we had one, and we did.

The AI enabler, and what it cannot do

An AI enabler is an AI worker like any other, and it makes most sense when a human enabler builds their own knowledge into it, so that it takes over part of their role. That person is usually also its manager. Like every AI worker, it needs a job description with a statable outcome. Its value rests on curated context about your platform and about what was decided and why. And it needs review, because the day it starts confidently misdescribing your setup is the day it becomes harmful.

It has one advantage over any human enabler, and it maps directly onto the story I opened with: asking costs nothing at all. No ledger, and no favour to repay. If the only problem were the social cost of asking, an AI enabler would close the gap completely.

It doesn't, for structural reasons.

It is not where beginners look first. For people taking their first steps in building AI, turning to an AI for help is not the most intuitive thing to do. They are still learning what these systems can do, and asking one for help with another is a step most of them haven't taken yet.

It cannot authorise. Much of being stuck is not a knowledge problem. Can I have access to that system? Is this allowed? An AI enabler can tell you what the rule says. It cannot carry your case or decide that in this instance the answer is yes.

It cannot notice that you have gone quiet. A system that waits to be asked cannot tell someone who succeeded from someone who tried twice and gave up. A human enabler notices the person who stopped coming, and following up with that person is often the most valuable thing the role does.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

It serves the people who cannot check it. Its users are, by definition, the least equipped people in the organisation on this subject, which is why they are asking. When it is confidently wrong about how your setup works, they have no way of noticing. Most AI workers are checked by someone with domain judgment. This one is checked by people who came to it because they don't have that judgment yet.

So the human enabler who manages it has to treat error review as the core of the job, and the human half of the function stays when the AI half arrives.

Two conditions that decide whether the role works

The first condition is that perceived availability beats allocated time.

This is where a protected share of someone's week usually fails, for a reason my opening story explains. If I sense that the person designated to help me is drowning, I go straight back to the private ledger. I ration. I bring the small questions and swallow the large ones, and the role's entire purpose is void while continuing to appear on the org chart and in the reporting.

What matters is how much time the people asking believe is there for them, more than how much you allocated. A protected twenty percent that people trust is worth more than an unprotected fifty percent they can see is under pressure. Practically: say out loud that this is the person's job, and be very careful about who else has a claim on that time, because the moment the day job wins twice in a row, everyone recalibrates, and recalibrating back is much harder than getting it right at the start.

The second condition is that the role is unconditional to ask, never unconditional to do.

The word I keep coming back to for this role is unconditional, and it needs a precise boundary or it destroys the thing it is meant to build. Unconditional applies to asking: any question, any number of times. It does not apply to doing.

In my six weeks, I did the work. He was there when I got stuck. That distinction is the whole difference between an enablement function and the way enablement functions usually end: as an order-taking queue. It happens gradually and for entirely reasonable-sounding motives. The enabler is faster at it, and the requester is busy. Six months later there is a request form and a backlog, and a set of managers who have AI workers built for them, running work they cannot explain, that they are nonetheless accountable for. You will have created exactly the unmanaged AI this guide warns against, and you will have created it through helpfulness.

The test is simple enough to apply in the moment: at the end of the interaction, who can do this next time?

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Take this with you: Ask the people who manage AI workers where they go when they get stuck. Write down every name you hear. If one name keeps coming up, that person is already doing the enabler's job without the title or the time for it. Make it official: say out loud that it is their job, and protect that time. If nobody gets named, people are getting stuck alone, and that is the enablement gap your check-up pointed to.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART III · CHAPTER 13

The context owner

I see it in my own work with AI every day. When I have kept my context in good shape, the outputs are excellent. When I haven't, I spend the conversation correcting: *yes, but that's no longer important.* The AI is working with exactly what I gave it.

It's no different with people. You can hire the best expert in the world, but if you don't give them access to your company's knowledge, or hand them outdated material full of contradictions, even they will struggle to produce anything meaningful. An AI worker is only as good as the context it works with.

Much of that failure comes from unclear responsibility. When nobody is clearly responsible for the context, nobody notices it drifting, and the AI worker pays for it.

The fix depends on what the context is for. When a piece of context is so specific that it serves one AI worker, that worker's manager can own it and manage it as part of the worker's anatomy. It gets harder when the context serves purposes beyond feeding an AI worker, or when it serves several AI workers. That is when the context itself needs a clear owner, because too many cooks spoil the broth. When everyone adds to a context file and changes it, a change that helps one user or one AI worker can cause serious problems for another, and soon nobody can say which version is true. Somebody has to own the truth and decide what is relevant and what goes in.

Organisations with good data governance have already learned much of this for structured data. Context governance takes it a step further: deciding which knowledge people and AI workers need to deliver, and how it stays relevant and current.

Who the context owner is

The context owner is usually the person you would naturally go to with questions about a project, a client or a field of expertise. They already hold that knowledge in their heads. The role gives them the say over the written version.

They decide what goes into the context and what doesn't, and they weigh which changes can be made without causing problems somewhere else.

It is therefore crucial that managers of AI workers and context owners, unless they are the same person, cooperate closely and keep an open flow of communication.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Why it pays off

One of the first things I do with a new client is curate the most relevant context. What are the most crucial things about the business that every AI worker should know? And then, area by area: the products, the brand, the clients, the rules. With the help of an AI worker built for curating content, this can be done quickly. But it has to be done deliberately if you want your AI to deliver real value instead of generic output. Once it exists, every AI worker starts from there.

And it builds something your competitors cannot buy. The model underneath is available to them on the same terms, but what your organisation knows about its clients and its own way of working is not.

Take this with you: Pick one AI worker and list the context it runs on. For each piece, write down who owns it. Where there is no name, ask who people naturally go to with questions about it, and make that person the owner. Then ask them one question: what in here is no longer true?

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART IV

Making It Work

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART IV · CHAPTER 14

The upside and the hard questions

At one of my first clients, there was a colleague who was very sceptical when I came on board. You could feel how worried she was about what all of this might mean for her job. Then one day she came to me with an idea for an AI worker. Goosebumps.

Of everything the method has done, that is the effect I had not dared to hope for.

The upside

That moment is where the story changes. "AI is coming for my job" becomes "I have an idea for an AI worker", and it happens when people get a way to shape what AI does in their work instead of waiting for it to happen to them.

The method also gives people control when everything around them is changing. Many are almost paralysed by the pace and the uncertainty. A method with simple steps, built on patterns that have worked before, gives them somewhere to start. Writing a first job description for an AI worker is already taking back control, one step at a time.

It makes room for conversations that never had a place. Once people talk about what an AI worker should do and who works with it, they start noticing what else is missing, such as a way to share what they learn or a place to talk about what isn't working. And they suggest how to close those gaps themselves.

It also gets you out of the tool debate. The method starts from the work that needs to get done, and then you build the AI worker with whatever your tool stack allows. Depending on your business, most AI workers can be built with what you already have, at least in a first version. The fully automated version may need an upgrade later. By then, though, the worker has already shown what it brings, and you have the data to argue that the extra functionality is worth its cost and its risk.

And it lets leaders steer the combined workforce strategically. When the fire drill comes, with its question of how many AI solutions are running and what value they deliver, the answer is already there.

The hard questions

"Isn't this just more bureaucracy?"

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Job descriptions for AI, when we don't even write them properly for people? A registry on top? It's a fair worry, and it becomes real the moment the process is over-engineered. If people spend hours writing job descriptions and running reviews, they won't do it. Full stop.

That is where AI itself helps. With my clients, we build a recruiter agent into their AI tool of choice, such as Copilot or Langdock. It asks the manager a series of questions, runs the risk assessment in the background against the company's own policies and guidelines, and creates the job description and the registry entry automatically. It takes minutes, not hours. Where a company already has processes for risk assessment or technical documentation, we link to them rather than reinvent them. A development lifecycle or an AI governance framework that is already in place is an advantage: the method builds on it.

Reviews need the same pragmatism. A simple, low-risk automation might get one review a year, and that review might be a single question: is this still needed? A complex or risky AI worker gets a closer look, more often. How you set this up depends on your processes and your culture, and there is no one size that fits all. What doesn't change is that it has to be easy for the people you ask to do it.

"Isn't this treating software like people?"

It's a fair point, and one I take seriously. The method borrows practices from managing people because they work, from job descriptions to reviews. What it doesn't borrow are names and faces. I'm careful about how AI workers are named, so nobody mixes them up with colleagues, and I have trouble with AI workers that come with human faces or names, because they make it easy to forget you are dealing with a machine. An AI worker is like a human colleague in many ways. The differences are the ones we should never forget.

"Doesn't the organisation now own my expertise?"

Possibly more of it than before, yes. If what you know ends up in context files, you give your organisation more than you might have done otherwise. But the question itself is older than AI: how much of your skills and knowledge belongs to your employer, and how much to you? AI only makes it more pressing. In the end it is a decision everyone has to make for themselves: does it feel fair to give my skills and knowledge to my employer?

"Isn't this a lot of work?"

Yes. You have to mean it. Done half-heartedly, the method won't show its benefits. It is simple and not overly complex, but it isn't magic. It needs people who are willing to put real thought into it, and to stay with it.

The good news is that you don't have to do all of it at once. The next chapter is about where to start.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Take this with you: Think of the most sceptical person on your team and ask them one question: if you could hire a colleague to take one task off your plate, which would you choose? Their answer might be your first AI worker.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART IV · CHAPTER 15

Starting small

The One-Workforce Method was built with an organisation-wide application in mind, but it was deliberately designed so that you can get started today. You can begin with your own team, within your own sphere of influence, because there is nothing you need from anyone else to get going. Where tools and guidelines already exist, you connect to them, sometimes with nothing more than a link or a note on where the information sits. Where nothing exists yet, the job description helps you decide which tools you actually need.

Check your clarity first

Before you think about AI at all, check whether your team knows where it is going. When the ground keeps shifting, even the best teams lose sight of their common north star: why they are here as a team, and where they want to go. The same is true for the values that guide the way. A simple test is to ask your team members. If they don't give you the same answer, you have groundwork to do first.

Write the first job description

The simplest step is a job description, and its first and most important question is who will manage the AI worker. One name, one person.

If you don't use AI yet, ask yourself where you could use another pair of hands if hiring budgets weren't an issue. The answers often point to work that AI can take over, or at least support.

If AI is already running in your team, go back and fill in a job description for each of your agents and automations. If you can answer the questions, you are already pretty far. If you can't, you know where to start.

You don't need a recruiter agent for this. The template on paper works. Or use it to guide a team discussion, record it, and give the transcript and the template to your AI tool to turn into a job description.

When someone asks why we bother

Someone will ask: "You want job descriptions for AI, when we don't even do them properly for people?" It's a fair question, and most organisations would have to admit the answer is no.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

Raise it yourself, and it becomes an invitation. The discipline you build for AI is the second chance for the human side. Build it once, and run it for both sides of the workforce.

Keep the list simple

You don't need an automated registry. For an AI workforce of three to five AI workers, a list you keep by hand in Excel or SharePoint does the job. What does matter is deciding early which metrics you want to see in your portfolio view. The job description is where you define what each one measures, so you can steer your AI workforce with them later.

Find your enablers

Look for the people on your team who already use AI and are eager to learn, and who see this as a chance to grow. Then give them protected time for the role, or it will stay a favour. They are your first enablers.

Check your own vital signs

As the leader, take the check-up in chapter 11. Look at where the warning signs are, and think through how you could close the gaps.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

PART IV · CHAPTER 16

Closing remarks

AI keeps evolving, and so does this method. I am implementing it with clients who let me try something new with them, and I learn with each one, because every business and every industry comes with its own needs and challenges. The method in this guide is the version I stand behind today. It won't be the last one, because environments change and people do too, and a method that doesn't flex to those changes will never change anything.

One thing has held with every client so far, and it is turning into a pattern: leaders have to mean it. Without that, the method won't work, however well the templates are filled in.

And in the end, the whole method is about people. That is my north star. If the method makes the people who work with it weaker instead of stronger, it has taken a wrong turn.

If you try it, I would love to hear what works for you and what doesn't. Get in touch. Every conversation makes the method better.

© 2026 Wawi Kantz e.U. · oneworkforcemethod.com

The guide continues online

The chapters live at oneworkforcemethod.com, together with the resources that go with them. New resources are added step by step and announced on LinkedIn.

oneworkforcemethod.com