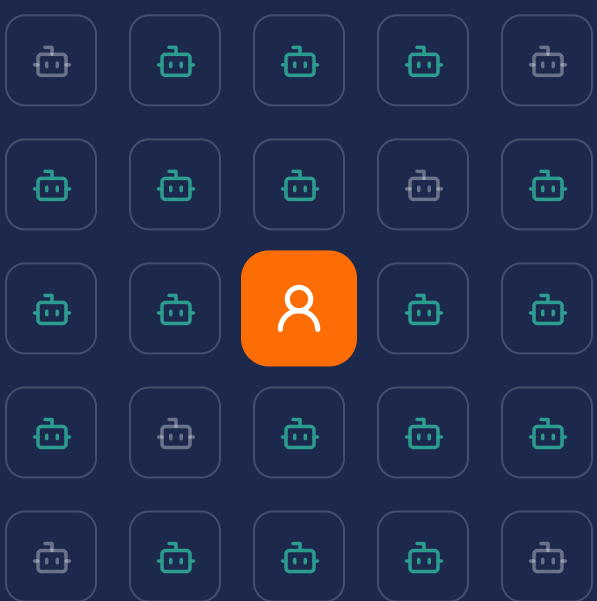
E-book | Free download

For professional services leaders

# The next chapter in the human workforce

Why AI will never replace us.





**Martijn van der Hoeden**
Founder & CEO of PSOhub

Published by
PSOHUB.

Contents

# Three parts, one argument.

The pessimistic case is real and I make it properly before I argue with it. Then what a bigger model does not touch. Then what to change on Monday if you run a firm.

**Part I** **The threat we are all being sold**
The pessimistic case, made properly. Layer by layer, from the Peak Collective to the C-suite.

**Part II** **What cannot be replaced**
Six things a bigger model does not touch, and the one that worries me most.

**Part III** **What this means if you run a firm**
Pricing, automation, data, apprenticeship, hiring, measurement. What to change on Monday.

— A note on where I stand

## I build the software that does this.

My company automates the operational spine of professional services firms. When someone tells me that AI will eliminate tasks across every level of an organization, my honest answer is yes. It will. It already is. This paper is about what happens next.

# 01

Part I

## The threat we are all being sold

The pessimistic case, made properly. Layer by layer, from the Peak Collective to the C-suite.

---

Open any feed today and the message is the same. Agentic AI is coming for the workforce. Not for a few repetitive roles at the edges, for all of it. Entire departments reduced to a prompt and a review step. The professional as an expensive relic.

I am not the right person to wave that away. I build the software that does this. My company automates the operational spine of professional services firms, and the whole point of an agentic roadmap is to remove work that people are currently paid to do. When someone tells me that AI will eliminate tasks across every level of an organization, my honest answer is yes. It will. It already is.

So let me make the pessimistic case properly before I argue with it. The disruption is not concentrated at the bottom of the organization the way previous automation waves were. It runs through the entire pyramid, and it hits each layer differently.




- **Roles merge.** Depth on demand, smaller executive teams.
- **Produces directly.** Plans and analyses no longer commissioned.
- **Squeezed both sides.** More output, less delegation.
- **Hit first.** The draft, the code, the write-up.
- **Touched last.** AI has no hands. Droids arrive later.

## The Peak Collective goes first

Young professionals are the first layer to feel it, and this surprises people who expected AI to start at the bottom.

Think about what the Peak Collective actually does in a law firm, an accountancy practice, a consultancy, or a development team. They write the first draft of the report. They produce the code. They run the intake interviews and write them up. They assemble the proposal. They draft the agreement from the template and the notes. They pull the numbers together so a senior can interpret them.

Every one of those tasks is now something an AI agent does faster, cheaper, and often better than a person two years into their career. Not equally well. Better, in the specific sense that matters to a partner reviewing it at eleven at night: fewer omissions, consistent structure, instant turnaround.

> **If a senior with the right tools can produce in an afternoon what previously took a junior a week, how many juniors does the firm actually need?**

That is a simple and uncomfortable question, and firms are already asking it in their planning meetings.

## The Mid-Sphere is squeezed from both sides

The senior accountant, the lawyer, the lead developer, the senior consultant. This layer looks safe and is not.

| Peak Collective | Mid-Sphere |
| --- | --- |
| Has adopted AI as a best friend. Constantly, casually, without asking permission. | Is struggling. The busiest people in the building, with authority built on fifteen years of knowledge. |

There is an adoption gap that nobody planned for and everyone can see. Not because these people are incapable, but because the pace of change right now is genuinely difficult to keep up with while also delivering client work.

Meanwhile management arrives with an AI-first strategy, and the instruction is effectively this: use the tools, do more yourself, stop delegating. So the senior takes on the drafting they used to hand down. The junior is no longer needed for that piece. The Mid-Sphere professional becomes more productive and more alone.

Something is lost in that trade that does not show up in any utilization report. Everything becomes output-driven. The journey is skipped in favor of the outcome. The work still gets delivered, and the process by which a person used to become senior quietly disappears.



## The management layer automates the layer below it

Managers are now doing work that used to be commissioned. A budget plan in minutes. A strategy document in seconds. A market analysis, a business case, a board summary, all produced directly rather than requested from the professionals below. The management layer is eating tasks upward from the Mid-Sphere, in exactly the same way the Mid-Sphere is eating tasks from the Peak Collective.

Here is where the pessimistic case reaches something more serious than headcount. The billable senior consultant is at risk, and that risk collides directly with the business model that most professional services organizations run on. Utilization. Billable hours. Bodies on a project.

> **If you no longer need the bodies, and the outcome is the same or better, why would a client keep paying for the hours?**

**The uncomfortable part**

**The threat is not that AI does the work badly. It is that AI does the work well enough to expose the fact that clients were buying time when what they wanted was a result.**

## C-level: the business model breaks before the jobs do

Which makes this the real C-level agenda, and it is not a technology agenda.

Hours cannot remain the revenue driver. The future has to be built on value-based billing, where the price is set by the outcome delivered rather than the effort consumed. That means packaging services into service-based products, with a fixed price and an agreed value, sold on what the client gets rather than on how long it takes you to get there.

I should be honest that the death of the billable hour has been announced regularly for twenty years and it keeps surviving. The difference now is that the cost of production is genuinely collapsing, and no pricing model survives a hundredfold change in the cost of the thing it is pricing.

There is a second effect at this level. The traditional executive team is a set of specialized knowledge holders: CTO, COO, CMO, CFO, each carrying depth that nobody else in the room has. When that depth is available on demand, I expect a number of those roles to merge. A CEO with good AI support can hold the CTO conversation credibly. The requirement shifts from having deep knowledge of every domain to knowing which questions to ask and which answers to trust. Smaller executive teams, broader mandates, more generalists.

I am aware that is a provocative forecast. I would rather put it in writing and be argued with than pretend it is not visible.

## Hands-on workers: not AI, but droids

The base of the pyramid is the layer AI touches last, because AI has no hands. The nurse, the electrician, the site foreman, the technician, the care worker. Their work is physical, situational, and built on human contact, and that combination is the hardest thing to automate. Human-to-human relationship is not a nice extra in this layer, it is the service itself.

That said, robotics arrives eventually. Heavy lifting, repetitive strain, toxic and dangerous environments, the work that damages the people who do it. Specialized droids will take that on, and once those machines are directed by capable AI rather than fixed programming, the range of tasks they can handle grows quickly. I am not sentimental about that. Work that breaks human bodies is not work worth protecting.

## The timeline: what arrives when

Artificial intelligence has been around for more than a decade, but the last three years are what changed the picture. Large language models get better every month, and generative AI has moved from a demonstration to something people use without thinking about it.

As with most innovations, it started with personal use and only then became accepted in the corporate workplace. Private use took off in 2022. Businesses are still struggling with how to regulate AI, and LLMs inside the corporate environment only arrived at scale in 2025.

The next phase is agentic. LLMs working in collaboration with business software through AI agents connected to CRM, PSA and ERP systems. General use of AI agents started at scale in 2026, and I expect general adoption to really flourish in 2027.

Letting those agents act autonomously in daily operations is a slower story, and it will vary by industry. Depending on the vertical, adoption will go faster or slower. My expectation is that autonomous AI is generally accepted by 2030.

2030 is also the year I expect superintelligence to be introduced, which changes the landscape again. That is the point where we enter an era in which no human can exceed AI in any field. Physical work runs on its own clock. Specialized robots and droids become common between 2027 and 2030, and by 2035 I expect them to be generally accepted for autonomous tasks, taking over a growing share of human labor.

**2022** Private use goes mainstream

**2025** Enterprise LLM use at scale

**2027** AI agents generally adopted

**2030** Autonomous AI and superintelligence

**2035** Droids at scale in physical work

## Where these dates come from

The band is my forecast, not a consensus view. These are the numbers I am reading it from, including the places where the published research argues against me.

**2022 — private use.** ChatGPT opened to the public on 30 November 2022, and consumer uptake ran years ahead of any corporate rollout. This is the only milestone in the band that is history rather than forecast.[1]

**2025 — the enterprise catches up.** McKinsey's 2025 State of AI survey, covering 1,993 respondents across roughly 105 countries, found 88 percent of organizations regularly using AI in at least one business function and 72 percent using generative AI, up from 33 percent a year earlier. The same survey found that nearly two thirds had not yet begun scaling AI across the enterprise. That is why I date arrival at scale to 2025 and not adoption at depth.[2]

**2027 — agents in general use.** McKinsey puts 23 percent of organizations at the scaling stage with an agentic system in at least one function, and another 39 percent experimenting. Gartner expects 40 percent of enterprise applications to carry task-specific agents by the end of 2026, up from under 5 percent in 2025, and a third of enterprise software to include agentic AI by 2028. The counterweight is documented just as clearly: Gartner also expects more than 40 percent of agentic projects to be canceled by the end of 2027 on cost, unclear value and weak risk controls. Wide availability by 2027 is well supported. Reliable production use is the harder half.[3]

**2030 — autonomous operation, and the contested one.** Gartner projects that at least 15 percent of day-to-day work decisions will be made autonomously through agentic AI by 2028, up from effectively zero in 2024, which makes broad acceptance by 2030 a reasonable extrapolation. Superintelligence by 2030 is the aggressive end of the range and I will say so plainly. Founder forecasts cluster around 2030, while surveys of AI researchers have put a 50 percent chance of human-level machine intelligence closer to 2047, and the Metaculus community median moved from 2055 in January 2022 to 2033 by mid-2025. My date sits with the builders, not with the surveys.[4]

**2035 — droids at scale.** Goldman Sachs Research has a base case of more than 250,000 humanoid shipments in 2030 and a $38 billion market by 2035, revised up sixfold from its earlier $6 billion estimate as component costs fell. Morgan Stanley models roughly 13 million humanoids in service by 2035, overwhelmingly in factories and warehouses, and expects adoption to stay slow until the mid-2030s before accelerating late in the decade. Common in industrial work by 2035 is well supported. Broad substitution for human labor sits later than my band implies.[5]

## So will we all be out of a job?

Add up the whole pyramid and the pessimists have a real case. Fewer juniors. Squeezed seniors. Merged executives. Automated physical work. A pricing model that stops making sense.

**I still do not believe it ends with us out of work.**

Something will change permanently in how we work, what we spend our days on, and which tasks we can honestly claim to do better than a machine. On a growing list of tasks, we cannot. I am not going to pretend otherwise in a paper that has my name on it.

But a task is not a person. We are human, and there is a specific set of things that makes us human which I do not believe gets replaced. Not by a bigger model, and not by a better agent. That is the rest of this paper.

# 02

Part II

# What cannot be replaced

Six things a bigger model does not touch, and the one that worries me most.

---

## We are sensing instruments, and nobody has built a replacement

Walk into a client meeting five minutes late and you know something within the first second. Someone has been crying, or arguing, or checking their phone under the table. The air is different. You adjust before you sit down.

None of that arrives as data. It arrives through skin, eyes, ears, and a nervous system that has been calibrating on other humans since birth. A model can read a transcript of that meeting and give you a competent summary. It cannot tell you that the CFO said “we’re aligned” while looking at the door.

We tend to describe the human senses as inputs, as if we were cameras with legs. That undersells it badly. Sensing goes beyond collection: the interpretation happens in the same instant. The consultant who notices that the client’s project manager has stopped saying “we” is doing forensic work with no visible evidence. That project is in trouble, and the dashboard will not know for another month.

Any business that involves being in a room with another human being runs on this. Which is most businesses.

## The sixth sense: knowing more than you can explain

There is a moment familiar to anyone who has priced a deal, hired a person, or approved a design. Everything on paper says yes. Something in you says no.

Sometimes you are wrong. Often you are not, and afterward you cannot fully reconstruct why. The philosopher Michael Polanyi called this tacit knowledge. It is the reason a master craftsman cannot write down the recipe, and the reason experienced people are worth what they cost.

Gut feeling gets treated as the soft opposite of data. It is the opposite of nothing. It is compressed experience, thousands of past cases pattern-matched below the level of language, arriving as a physical sensation rather than a sentence. Your body reaches the conclusion before your vocabulary catches up.

This is the part of us that looks most like AI and is least replaceable by it. Both are pattern recognition at scale. The difference is what the patterns are made of. A model is trained on the record of what happened. A person is trained on having been there when it happened, with something at stake, and with the memory of how it felt when it went wrong.

**You cannot train that on someone else's data. You have to live it.**

## Nothing is at stake for the machine

Here is what I keep coming back to, and it is less about capability than about consciousness.

An AI can produce a recommendation. It cannot care whether the recommendation is right. It will not lie awake over the redundancy it advised. It has no reputation to lose, no family depending on the outcome, no colleague whose trust it will have to face on Monday. It does not experience the consequence of its own judgment.

> **Judgment without exposure is not judgment, it is output.**

A bigger model will not close that gap, because it is a difference of category. Every serious human decision carries an author. Someone who signs, who is answerable, who can be asked "why did you do that?" and has to answer in front of people who will remember. Accountability is what makes it a decision at all. Paperwork attached to a decision after the fact cannot supply it.

I have never met a client who wanted the machine to sign off. They want a person who is willing to be wrong in public and stand behind it anyway.

## The evidence on sensing

I am not asking anyone to take this on faith. Three separate bodies of research point the same way: the market pays for social perception, employers say they need more of it, and it measurably predicts whether a group performs.

**The labor market already prices it.** David Deming's study of two national longitudinal cohorts found the wage return to social skills is positive even after controlling for cognitive and non-cognitive skill, that cognitive and social skill work as complements rather than substitutes, and that the complementarity has grown over time. Jobs requiring high levels of social interaction grew their share of employment while less social work shrank. A one standard deviation increase in social skills was worth $2.67 an hour against a baseline of $16.23.[6]

**Employers are saying it out loud.** In the World Economic Forum's Future of Jobs Report 2025, analytical thinking is the top core skill, followed immediately by resilience, flexibility and agility, and by leadership and social influence. Empathy and active listening sit in the top ten. Leadership and social influence rose 22 percentage points as a named core skill against the 2023 edition, the largest increase in the survey, while manual dexterity, endurance and precision show the steepest declines.[7]

**Reading people predicts what a group produces.** The Science study that established a measurable collective intelligence factor found the strongest predictor was not the average or the maximum intelligence of the members, but their average social perceptiveness, measured by a test of reading emotional state from the eyes alone. A follow-up found the same effect in online groups with no visual cues at all, which suggests the underlying skill is inferring another mind rather than watching a face. A later replication attributed far more of the variance to individual intelligence, so treat the size of the effect as contested and the direction as robust.[8]

> **"We know more than we can tell."**
>
> **Michael Polanyi · philosopher**

## The work is the relationship

In professional services, firms like to believe they sell expertise. They mostly sell trust.

Clients do not renew because the deliverable was technically excellent. They renew because someone called them before the bad news broke. Because a difficult conversation about scope was handled without anybody losing face. Because they felt understood at a moment when they were under pressure.

Interaction is not the packaging around the value. In a large share of professional work it is the value. Persuading a room that disagrees with you. Telling a client something they do not want to hear while keeping the relationship. Getting two departments who dislike each other to commit to the same deadline. Noticing that a talented colleague is quietly burning out.

AI is genuinely useful in the space around all this. It prepares you, briefs you, remembers everything, drafts the awkward email you then rewrite. It does not do the moment itself. The moment requires two nervous systems in the same room, both of which can be hurt.

## Experience has to be earned, and this is where I am actually worried

This is where Part I comes back, because the squeeze on the Peak Collective has a consequence that no workforce plan currently accounts for.

Experience compounds, but only if you get to have it. Judgment comes from doing work badly, being corrected, and doing it again. The junior consultant learns by producing the first draft nobody keeps. The junior developer learns by fixing the bug they created. That work is unglamorous, inefficient, and exactly the work AI now does in seconds.

If we automate the entire bottom rung, we should ask honestly where the next generation of Mid-Sphere professionals comes from. You cannot skip to intuition. There is no shortcut to a gut feeling, and a career built entirely on reviewing machine output does not produce the tacit knowledge that made the previous generation worth hiring.

> **The sharpest risk in this paper**
>
> **It is not that AI replaces experienced people. It is that we quietly stop producing them.**

When the journey gets skipped in favor of the outcome, we are consuming a stock of human judgment that we are no longer replenishing.

That is a design problem and it is solvable, but only deliberately. Apprenticeship has to be rebuilt on purpose, because the old version of it was a side effect of inefficiency we are now busy removing.

## New work appears, and it always looks like nothing at first

Every wave of automation has produced this argument, and every time the pessimists have been directionally wrong about employment and roughly right about disruption. Both things are true, which is why the debate never resolves.

> **“Automation kills tasks, not occupations. The occupation reorganizes around whatever the machine cannot do.”**

Some of what is emerging is already visible. People who design and supervise agent workflows. People accountable for what an autonomous system did at three in the morning when nobody was watching. People who audit AI decisions for bias and defensibility. People whose entire value is judgment at the point where automated output meets a real client.

Most of the new roles are not visible yet, which is normal. In 1995 nobody had cloud security on a business card.

What I am confident about is the shape of the shift, not the job titles. Less production, more direction. Less doing the work, more deciding what work is worth doing and standing behind the result.

# 03

Part III

## What this means if you run a firm

Six moves, and only one of them is a technology project.

### 01 Fix the pricing model before it breaks you

If your revenue is hours and your cost of production is collapsing, you are on the wrong side of your own efficiency. Move to value-based billing and service-based products with a fixed price and an agreed outcome. Firms that make this move early will sell the AI leverage as margin. Firms that wait will hand it to clients as a discount they did not choose to give.

### 02 Automate the friction, not the humans

Every hour taken out of admin is an hour that can go to a client, a colleague, or a decision. If the hour goes nowhere, you have cut cost and lost capability. That is a bad trade dressed as a good one.

## 03 Unify your data before you get excited about agents

This is the least romantic point in this paper and the one I would defend hardest. An agent working across contracts, projects, time, and invoices that live in five disconnected systems will confidently produce nonsense. Unified data is not an AI feature. It is the precondition for AI being anything other than a demo.

---

## 04 Rebuild apprenticeship on purpose

Decide which work the Peak Collective keeps doing even though a machine could do it faster, and be explicit that you are buying future judgment with present inefficiency. Also close the adoption gap in the Mid-Sphere, because a senior layer that cannot use the tools will be managed around rather than developed.

## 05 Track time and AI tokens together

Value-based billing only works if you know what the outcome cost you to produce, and that now takes two meters instead of one: the hours your people spend and the tokens your agents burn. Measure both and you can see margin per deliverable rather than guess at it. There is an upside beyond cost control. Once your agents and models are trained on your own knowledge, the decades of judgment sitting in your files stop being overhead and become something you can package, license and sell. That is how you commercialize experience you have already paid for.

## 06 Hire for the things that do not automate

Perception. Nerve. The willingness to be accountable. These have always been undervalued in hiring because they are hard to test for. Their price is about to go up.

Closing

## The combination that wins

**I do not think we are heading into a smaller human workforce. I think we are heading into a stronger one.**

The parts of work that felt most modern, the processing, the reporting, the drafting, the reconciling, are exactly the parts a machine now does faster, more accurately and more consistently than we ever managed. That is good news. It clears the desk. What is left is older than any of it. Reading a room. Trusting a hunch. Sitting with someone through a decision that could go wrong. Being the one who signs. A machine can mimic all of that convincingly. It cannot mean it, and clients can tell the difference.

So the advantage comes from combining automation and craftsmanship on purpose: AI carrying the production at a speed and accuracy no team can match, people carrying the sensing, the judgment and the willingness to stand behind the result. Firms that build that combination deliberately will be more useful to their clients than they have ever been, and better paid for it. That is the silver bullet, and it is available to anyone prepared to design for it rather than wait for it.

# That is the next chapter.

Not humans against machines, and not humans as machine operators. The two working as one, with people doing the part that was always the point, finally freed from the paperwork around it.

**Disclaimer**

Written in the summer of 2026, and dated by design. The timeline and the forecasts are my own judgment rather than consensus, and in a field moving this fast some of them will age badly. Nothing here is investment, legal, tax or professional advice. My interest is not neutral and I would rather say so than imply it: I build and invest in software that automates professional services work. Research attributed to McKinsey, Gartner, the World Economic Forum, Goldman Sachs, Morgan Stanley and the studies cited belongs to those sources and is quoted as published; any error in summarizing it is mine. The views are mine alone, not those of any company, client or board I am associated with.

**About the author**



**Martijn van der Hoeden**
Founder & CEO of PSOhub

"I have built software companies for thirty years, I am active in Tech M&A, I host a podcast, and I have had this conversation with a lot of people in the field. I am also still invested in Tech companies. The world is changing fast, and behind all the threats I see enormous opportunity. As always, the winners will be the ones who adopt the change rather than fight it."

---

## Sources

1 OpenAI, "Introducing ChatGPT," 30 November 2022.

2 McKinsey & Company, "The state of AI in 2025: Agents, innovation, and transformation," November 2025.

3 McKinsey & Company, "The state of AI in 2025," November 2025 (agentic scaling and experimentation rates); Gartner, "40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026," 26 August 2025; Gartner, "Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," 25 June 2025.

4 Gartner, "Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," 25 June 2025 (15 percent of day-to-day work decisions made autonomously by 2028); Grace et al., AI Impacts expert survey, 2023 wave (n = 1,714), 50 percent probability of high-level machine intelligence by 2047; Metaculus community forecast for the first general AI system, 2055 in January 2022 versus 2033 in June 2025.

5 Goldman Sachs Research, “The global market for humanoid robots could reach $38 billion by 2035”; Morgan Stanley Research, “Humanoid Robot Market Expected to Reach $5 Trillion by 2050.”

6 David J. Deming, “The Growing Importance of Social Skills in the Labor Market,” Quarterly Journal of Economics 132, no. 4 (2017): 1593–1640 (NBER Working Paper 21473).

7 World Economic Forum, The Future of Jobs Report 2025, Chapter 3: Skills Outlook, January 2025.

8 Anita Williams Woolley, Christopher F. Chabris, Alex Pentland, Nada Hashmi and Thomas W. Malone, “Evidence for a Collective Intelligence Factor in the Performance of Human Groups,” Science 330 (2010): 686–688; David Engel et al., “Reading the Mind in the Eyes or Reading between the Lines?” PLOS ONE 9 (2014): e115212; for the contrary finding, Timothy C. Bates and Shivani Gupta, “Smart groups of smart people,” Intelligence 60 (2017).