

[THE PEOPLE PILLAR]

Your team isn't the
audience for AI.

They're the
operating system.

PAPER 2 OF 3 IN OUR AI-READINESS SERIES



This is not the soft paper.

Most readers will see the title *The People Pillar* and assume that this paper is about training, communities of practice, change management. It isn't.

This paper argues that the hardest part of becoming AI-ready isn't the tools or the strategy; it's the human system underneath. The org chart, the incentives, the fear, the identity people have built around the work AI is now doing.

Most companies are trying to install AI into a workforce that has not been told what its job becomes on the other side. It's also the pillar most likely to be skipped, because it's the hardest to measure and the slowest to move.

Yet, it's also most likely to kill an AI transformation when it is ignored. Without it, your technology investment sits unused and your strategy stays on a slide.

This is the paper about that gap. Read it with your own org chart in mind. That's where it stops being theory.



SEAN KING AND SVEN SABAS
CO-FOUNDERS, DRAGONFLY

A note on order: The Technology Paper came first as the most measurable pillar. We're publishing People before Strategy because, across the hundreds of stacks we have mapped, neither of the other two lands without it.

01	The reality check	05
02	Reframing people	06
03	The Three Stances	09
04	Where do your roles sit?	14
05	Measuring what matters	16
06	Culture as pre-condition	20
07	Where Dragonfly comes in	21
08	What's next?	23

Everyone's using it. The work looks the same.

Pick any company that announced an AI rollout in the last twelve months. The press release reads like every other one. The reality on the ground is harder to read and it isn't the story most people expect.

The expected story is that nobody's using the tools. Sometimes that's true. But it depends entirely on which kind of tool you mean, and lumping them together is the first mistake.

The AI add-ons bolted onto existing software (Notion AI, the Slack assistant, the copilot inside your CRM) often do follow the dead-on-arrival pattern: a spike of curiosity in week one, a plateau by week four, a quiet majority who never went back. For these, low usage is real.

But the core tools (ChatGPT, Claude, Gemini) usually tell the opposite story. Usage is high. It's sticky. People open them every day and would be annoyed to lose them. By every dashboard a vendor can show you, this looks like a runaway success.

And yet six months in, the executive team is still asking the same question: what has this actually changed?

Output looks similar. Cycle times look similar. The org chart is identical. The work, looked at closely, is being done the way it always was, just with a chatbot open in another tab.

This is the part most companies miss. Usage is not adoption. A team can use a tool intensively every day without changing how they work in any way that matters. The licences are justified, the dashboards are green, and nothing underneath has moved. In the adoption data we see at Dragonfly, this is by far the most common shape.

So the real test isn't how many people are logging in. It's this: if you removed the tool tomorrow, would the work change or would people be mildly slower for a week and then carry on exactly as before?

If the honest answer is "slower for a week," you don't have an adoption success. You have an expensive way of doing the old work.

The instinct, when leaders sense this, is to treat it as a people problem: more training, more incentives, a more intuitive tool, hire a Head of AI Transformation. None of these touch it. The work that matters is somewhere else.

“

Usage is not adoption. A team can use a tool every day and change nothing. That's not success. It's an expensive way of doing old work.

Your people aren't AI's audience. They are its operating system.

In our Technology paper, the first pillar of the Dragonfly framework, we reframed software. We argued that your tech stack isn't just a set of tools, it's the implementation of a process. If you don't understand the process underneath, the tools on top can't help you.

People work the same way.

Most companies treat their workforce as the audience for AI. The thing AI is being delivered to. The user base for the tool. So they communicate, train, support, and measure adoption. They run the AI rollout the same way they'd run any other software rollout, with the workforce as a recipient.

This is the wrong frame.

Your people are not the audience for AI. They are the operating system it runs on. Every workflow in your business is implemented through humans, their judgement, their relationships, their institutional knowledge, their incentives, the rituals they've built around their work. AI doesn't sit on top of this layer. It rewires it.

You cannot install AI into this system without changing the system. And changing the system is not a training exercise. It's a redesign of how your business asks people to spend their time, what it rewards them for, and what it tells them their job actually is.

“

It was never that your team couldn't use AI. It's that nobody has answered the question they're quietly asking... what does my job become on the other side of this?

The three layers no one is looking at.

01

IDENTITY



Your team built careers around tasks AI now does.

- The marketing manager who built their career around their ability to write briefs that made the agency relationship work.
- The senior engineer who became senior because they could read a 10,000-line codebase and hold it in their head.
- The customer success lead who built their reputation on the quality of their quarterly business reviews.


The work that defined them, the thing they got promoted for, the thing their identity is wrapped around, is now something a model can do in thirty seconds.

This isn't fear of change.

It's a rational response to an unanswered question. If the work that made me valuable is now automated, what makes me valuable?

02

INCENTIVE



Roles, titles, KPIs, and team structures were designed for a workflow AI is dismantling.

A customer success lead is measured on tickets resolved. If AI resolves the tickets, the metric is incoherent and the individual is left wondering what they are being measured on instead.

You can train someone on prompts all day. If their job description still rewards them for the work AI now does, the training is just theatre.

03 LEADERSHIP



Adoption is what leadership can see. Identity threat and org chart incoherence are not. So when usage dashboards stay flat, the conclusion is that the workforce isn't engaging.

The workforce is engaging. They are engaging with the question leadership hasn't answered. "What does my job become on the other side of this?" Until leadership answers it explicitly, in language that includes the word "role" and not just the word "tool", the dashboards will stay flat, no matter how many training sessions are run.

Your team isn't resisting AI. They're resisting an unanswered question about their future.

🚩 THE BOTTOM LINE

Tools change in weeks. Strategies change in months. Human systems have traditionally changed in years, and only once the question your team is asking has been answered honestly. Until then, every AI investment will stall short of the return that was promised. Not because the tools are bad, but because the operating system underneath isn't ready to run them.

Permitted. Embedded. Reconfigured.

The instinct is to wait for a standard way of working before rolling AI out widely. But the standard way is discovered through experimentation, not designed in advance. Waiting for one produces neither.

Across the companies we've worked with, that redesign sits at one of three observable stances: *Permitted*, *Embedded*, *Reconfigured*.

Most are *Permitted*. A few are *Embedded*. Almost none are *Reconfigured*, and that shape is itself part of the story.

Each stance is defined by what has structurally changed about how the work is done, not by sentiment, intent, or the language in the all-hands.

Stance is observable.

It's measurable.

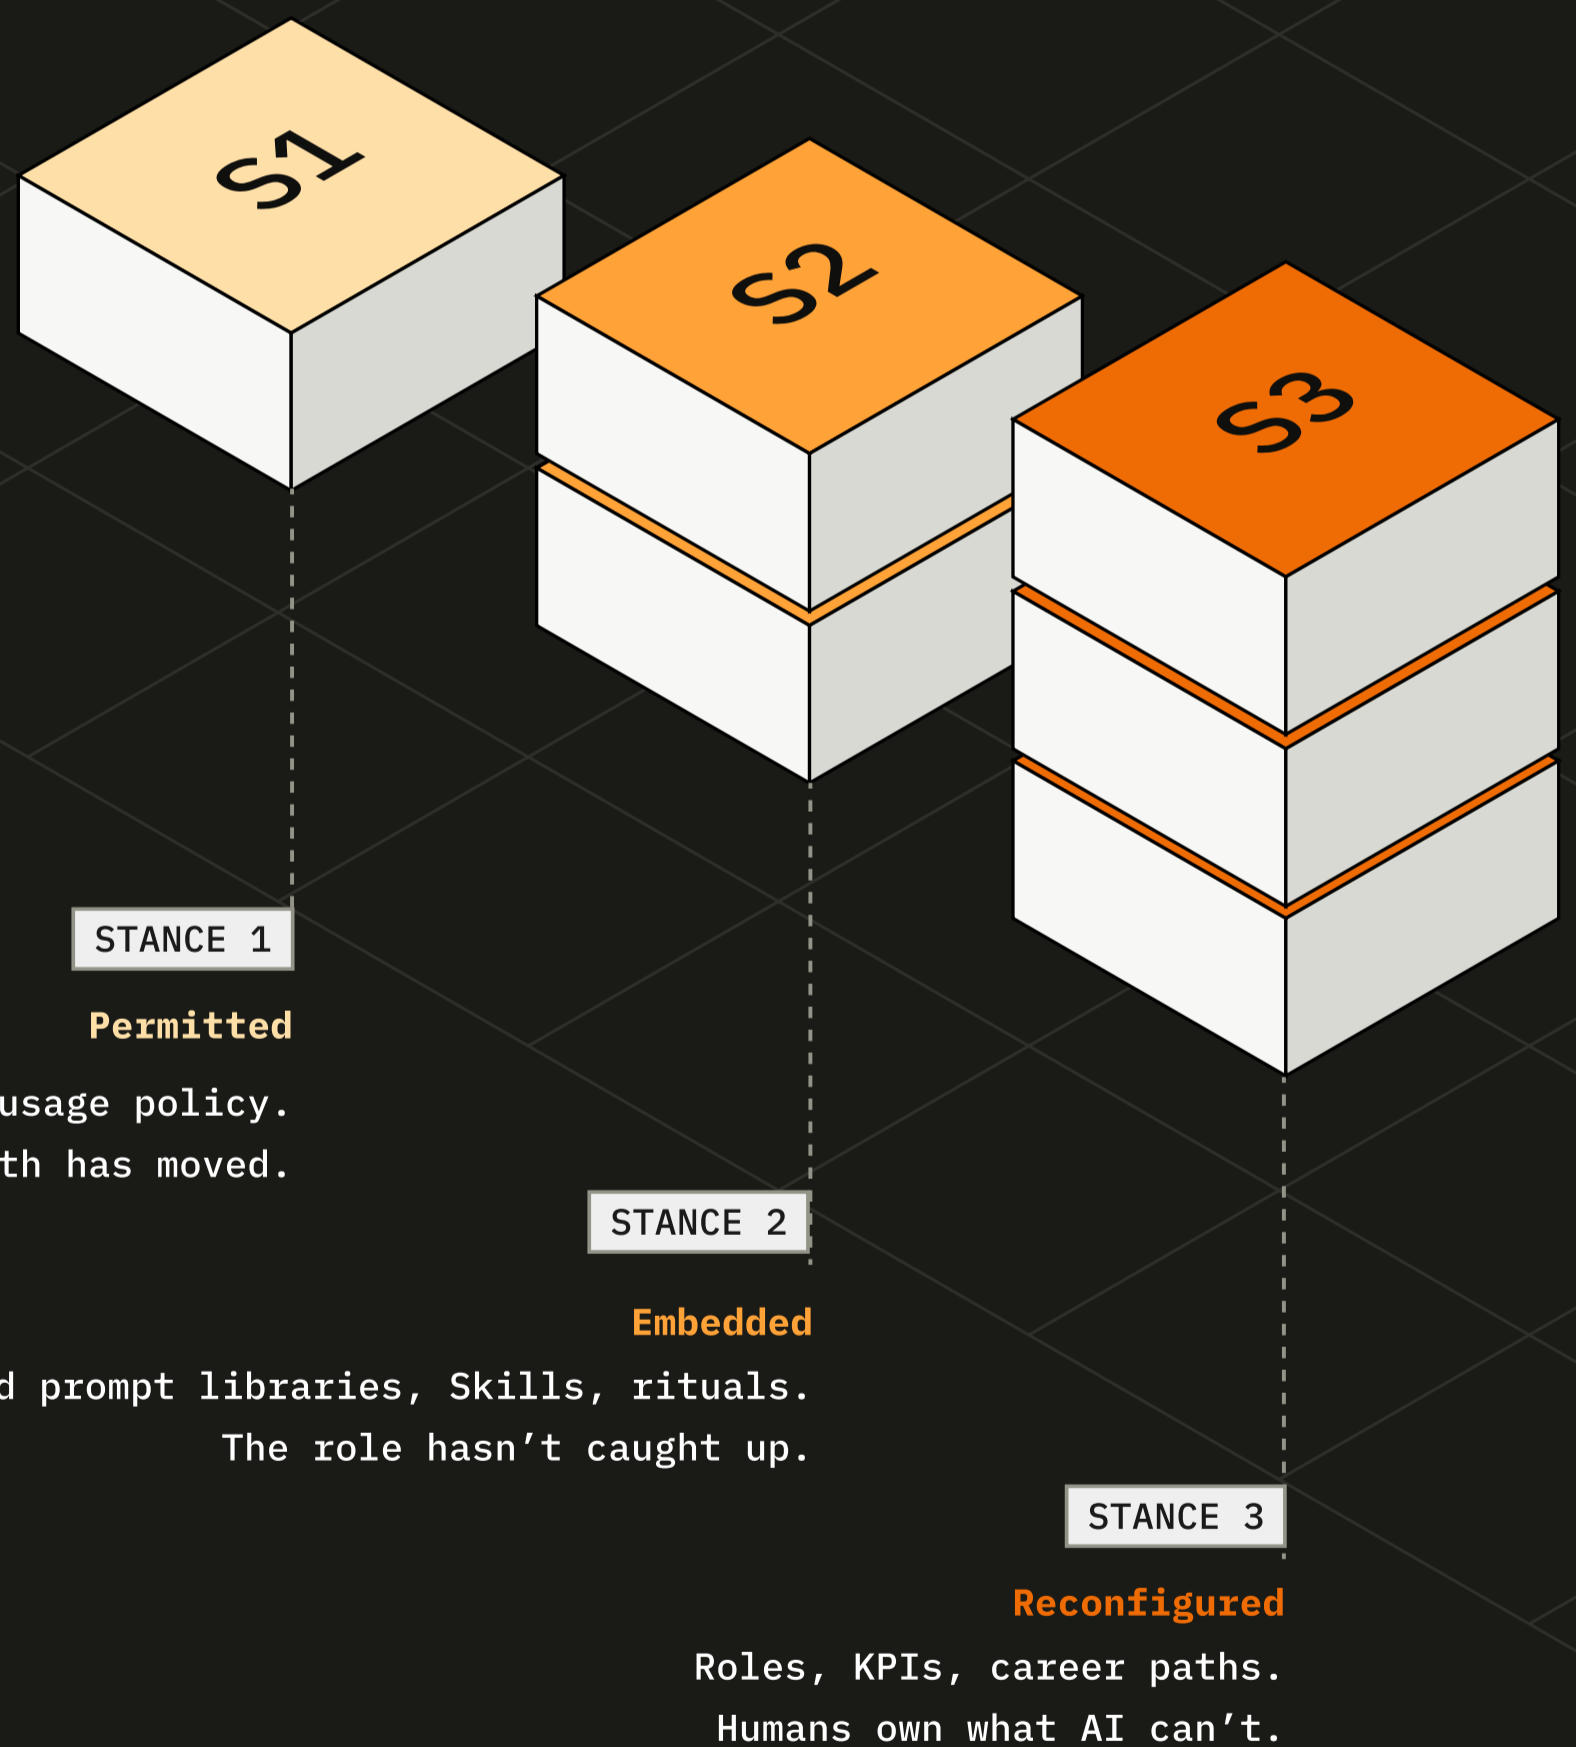
And it's almost always different from where leadership thinks the company sits.

Here's what each stance looks like, and the tell that gives away which one is really yours.

Tools, licences, a usage policy.
Nothing underneath has moved.

Shared prompt libraries, Skills, rituals.
The role hasn't caught up.

Roles, KPIs, career paths.
Humans own what AI can't.



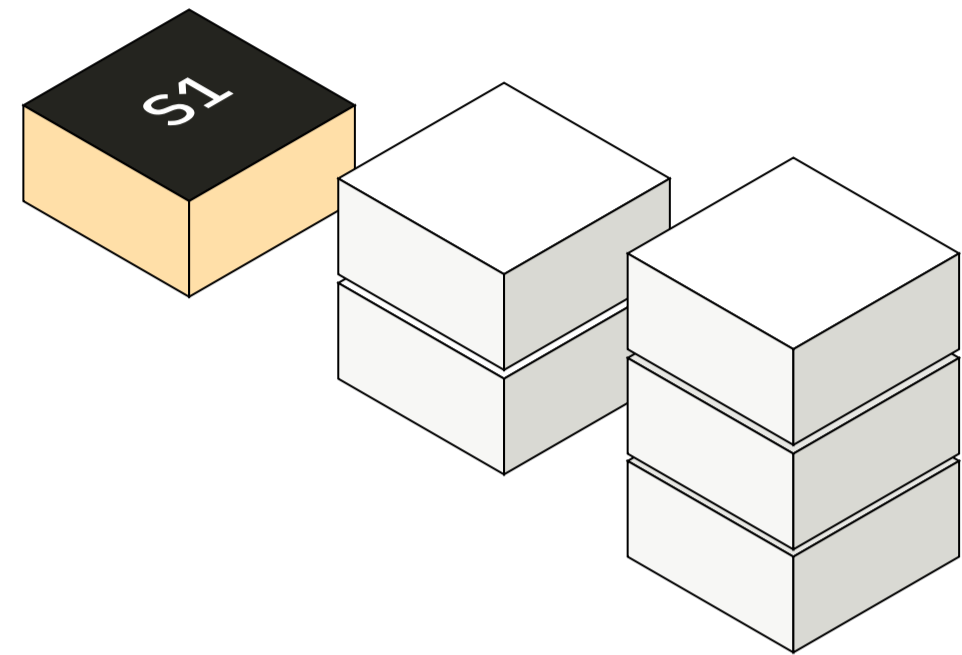
STANCE 1

Permitted

AI is allowed. Nothing about how the work happens has actually changed.

Tools have been rolled out. A licence policy exists. Some sort of AI usage guideline has been emailed around. A handful of power users have built genuine practice; most of the workforce uses AI occasionally for low-stakes tasks: drafting an email, summarising a meeting, asking a question they could have Googled.

Crucially, the org chart, the KPIs, the rituals, and the review cycles are unchanged. Job descriptions reference AI with vague aspirational language, if at all. Managers have no real visibility into who is using what or how well.



Leadership reports adoption based on licences allocated, not work transformed.

This is the dominant stance we see today. It's also the one most likely to be misread because on the surface, tool usage, prompt volume make green dashboards.

THE TELL

If you removed AI tomorrow, would anyone's job be visibly different next week? At Stance One, the answer is no. Slightly slower, perhaps. But not different.

STANCE 2


Embedded

Teams are deliberately working with AI. The job hasn't been redesigned, but the practice has. Adoption is no longer accidental.

Teams have shared prompt libraries, shared Claude Skills, internal communities of practice, and rituals that explicitly include AI: "draft it in Claude before peer review," "generate three options in ChatGPT before picking," "run the post-mortem through Claude." Managers know who their power users are, and have started to surface their patterns to the rest of the team.

This is where the culture you set earlier starts to show its hand. Champions either get championed: given air cover, visibility, time to share what they're learning, or they get held back by red tape, security reviews that never close, and policies written for a world that no longer exists.

Same people, same tools, completely different outcomes depending on what the organisation actually rewards.



Leadership has begun asking the harder question. "What does this role become?" appears in offsite agendas. Some teams have started informal experiments with role redesign: what if we removed this part of the job, what if we expanded the scope here.

The work has visibly changed. The role definition is catching up but hasn't yet been formalised.

THE TELL

This is a transitional stance. Companies don't stay here long; they either institutionalise the changes (Stance Three) or lose them (regress to Stance One) when the original advocates leave or the next priority lands.

STANCE 3


Reconfigured

The job has been rewritten and the human role is defined around what AI can't do.

Job descriptions, KPIs, team structures, hiring patterns, and career paths reflect a workforce in which AI is a teammate, not a tool. Humans are explicitly positioned around the work AI cannot do well: judgement under uncertainty, relationships, accountability, taste, edge cases, and the design of the systems themselves.

The champions who got air cover at Stance Two are now the people writing the new playbooks. The rituals that lived in pockets of the org have been institutionalised; they show up in onboarding, in performance reviews, in how the company hires.

The headcount picture has shifted. Some roles are gone or radically smaller. Some roles are new and didn't exist eighteen months ago.



The remaining roles look meaningfully different from their pre-AI versions, and the people in them can describe that difference clearly when asked.

Most importantly, the question “what does my job become?” has a clear, confident answer. And the answer is more interesting than the old job, not a diminished version of it. This is the difference between a workforce that fears AI and one that has been given a future to walk into.

THE TELL

This is a transitional stance. Companies don't stay here long; they either institutionalise the changes (Stance Three) or lose them (regress to Stance One) when the original advocates leave or the next priority lands.

The questions leaders are actually asking

Underneath the reframe sit a handful of intensely practical questions every leader rolling out AI is wrestling with right now.

They rarely make the whitepapers, but they're what the Monday meeting is about:

- Should everyone get the same tools, or different tools for different roles?
- Should we standardise on one model, or let people choose between ChatGPT, Claude, and Gemini?
- Should we be buying more tools, or consolidating to fewer?
- And underneath all of them: what does "good adoption" actually look like for us?

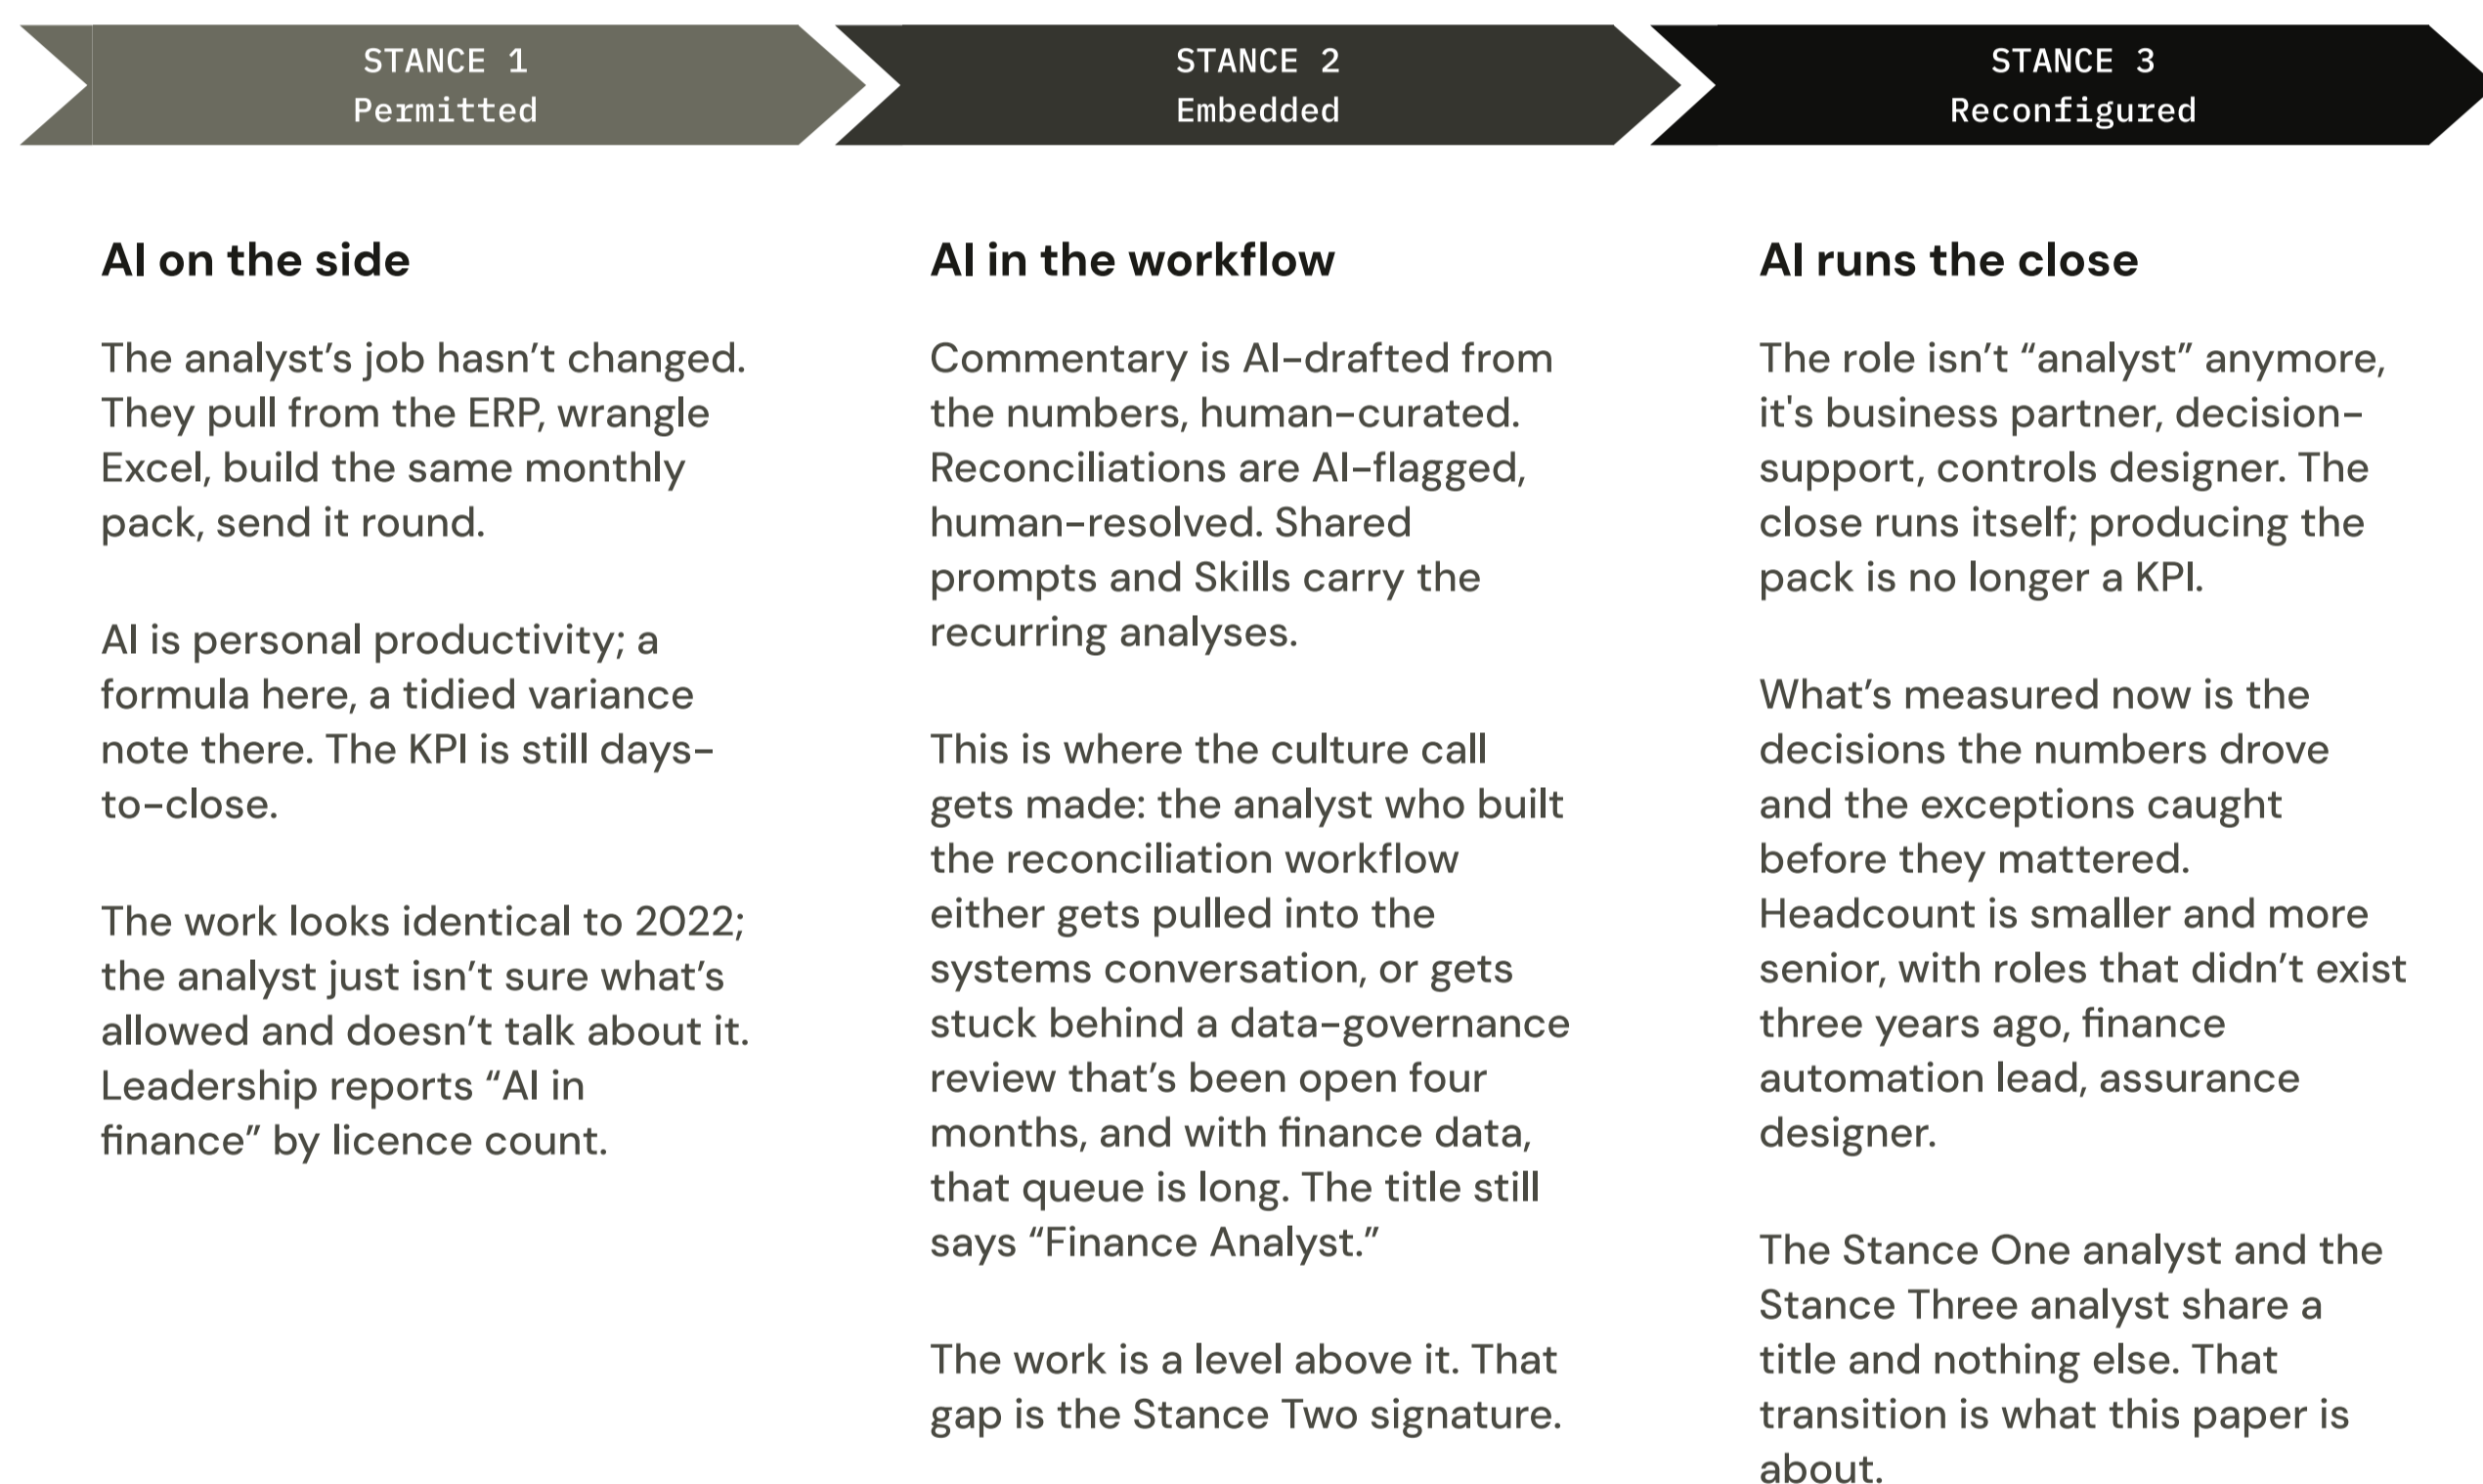
The uncomfortable part: none of these has a universal answer. The right tool spread, the right model policy, the right number of tools: all of it depends on how your people actually work, which is the one thing most leaders can't see.

The companies that get these calls right aren't the ones with the strongest opinions. They're the ones who can look at what's happening across their teams and let the pattern answer the question. Giving leaders that view is the problem Dragonfly exists to solve; we come back to it at the end of this paper.

Where do your roles sit?

As we did with Sales and Support in the first paper, we've mapped two functions across the three stances. These are illustrative, not prescriptive. The point isn't to copy the chain. It's to recognise the shape.

Finance: The Reporting Chain




The gap between Permitted and Reconfigured isn't just about using more AI. It's a fundamentally different relationship between your people and their work.

Operations / Business Systems: The Process Chain

Operations is where this is most live, and most uneven.

A team at Stance Three inside a company at Stance One is common, and so is the reverse. The friction back-office leaders are living with right now is that gap. And it doesn't close with better tools or more training. It closes when the company answers the question its workforce is already asking.



AI on the side

The AI features inside the existing tools get switched on; the ERP assistant, the service-desk copilot.

The team summarises a ticket, drafts an SOP. The process runs as it always has: handoffs between systems that don't integrate, exceptions cleared by hand, the same swivel-chair work.

The metric is throughput. The work looks faster in places, but not different.

AI in the workflow

AI is in the process, not just the tools. Exceptions are AI-triaged and routed with a recommended fix.

The system owner has started stitching agents across the platforms that never integrated; the thing they always wanted budget for.

This is where the culture call shows up sharpest: the lead rebuilding order-to-cash around agents either gets air cover, or gets told to wait while governance and procurement catch up.

Senior people move from pushing transactions to designing exceptions. The team is more productive and the career ladder is starting to look incoherent.

AI runs the close

"Business systems" has split. A few senior people orchestrate the agents running the transactional flow; humans set the rules, own the exceptions, design the controls.

The KPI is no longer throughput, it's process integrity, and the exceptions prevented before they cost anything.

The pyramid is flatter: fewer transactional roles, more systems-design and assurance, a new band of process orchestration leads.

Hiring shifts from "can you run the process" to "can you govern the agents and own the outcome."

The gap between Permitted and Reconfigured isn't more automation. It's whether your people run the process, or redesign it.

You're measuring activity. Not real change.


If your AI metrics still look fine in a world where nothing has actually changed, you're measuring the wrong things.

Most companies measuring AI adoption today are measuring activity, not change. Licences allocated. Prompts run. Hours "saved", almost always self-reported, almost always inflated.

These metrics have one property in common: they look identical whether AI is transforming the business or sitting unused.


The right metrics depend on who's asking. A CEO, a department head, and an individual contributor are measuring different things, and conflating them is part of why the picture gets muddled.

Below are the questions each audience should be answering, and the metrics that come closest to answering them today.



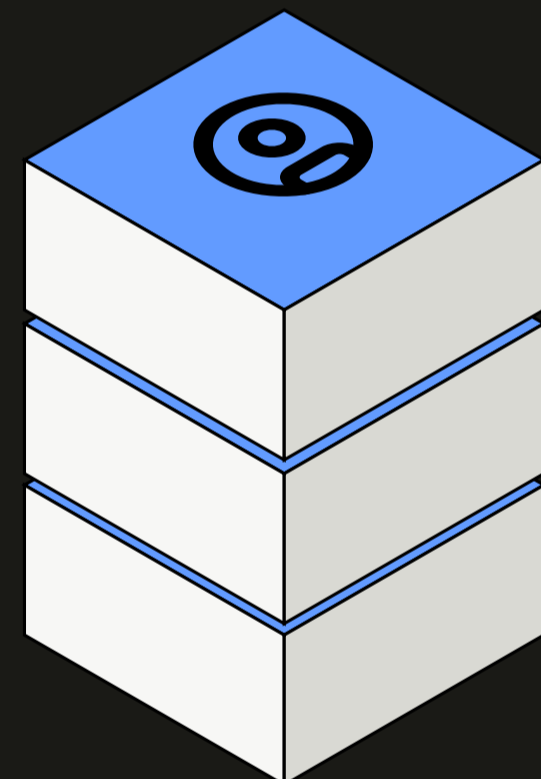
[BUSINESS]

Is AI changing what the company can do?



[TEAM / MANAGER]

Is the work actually changing?



[INDIVIDUAL]

Is my job actually different?

Note: the questions are the framework. The metrics are illustrative. The right specific KPIs will depend on your business, but the diagnostic questions don't change. These are also the questions Dragonfly uses to benchmark companies against their peers.

Is AI changing what the company can do?



This is the CEO's question, and it isn't "are people using AI." It's whether the company is materially different; in capacity, in economics, in what it's capable of producing compared to what it was before AI was deployed. If the answer is no, then whatever else is happening, transformation isn't.

There's a sharper version of this question for boards: are we accumulating champions or losing them?

At Stance One this metric doesn't exist. At Stance Two it's the leading indicator of whether the company will make it to Stance Three.

The companies that institutionalise the practice and the companies that regress to Stance One look almost identical on a Copilot dashboard, and completely different on the org chart eighteen months later.

📊 BUSINESS METRICS

<p>Output per employee, trended against headcount</p> <p>Revenue, units shipped, tickets resolved, content produced. The honest version of "AI ROI."</p>	<p>Revenue per AI-touched workflow</p> <p>For processes that have been redesigned, is the business outcome better, not just faster?</p>	<p>Cost-to-serve and cost-to-deliver</p> <p>Has the unit economics of how the business operates actually shifted, or just the line items?</p>
<p>Headcount composition</p> <p>Not "did we cut staff," but "are we hiring different roles than two years ago?" If new role types aren't in your hiring plan, you're not where you think you are.</p>	<p>Cycle time on core processes</p> <p>Time from intent to outcome on the workflows AI touches. The through-line back to paper one: processes are still the unit of analysis.</p>	<p>Champion retention and reach</p> <p>Are the people who built the early practice still here and leading wider remits? Or have they quietly left for companies that gave them air cover?</p>

✍️ DIAGNOSTIC TEST

Can you point to a P&L line that has moved because of AI, or are you measuring activity? And as a second test: can you name the three people in your company whose work has most visibly changed in the last year, and tell me whether they got promoted, ignored, or have one foot out the door?

Is AI actually changing the work?



The team lead's question is closer to the ground. Not "do we have the tool" but "is the work being done differently, and is it better."

This is where most companies have the worst visibility, because the data lives in tool-level usage logs that don't aggregate up to anything useful.

The team is also where the Stance Two signature shows up first. The manager who can name their power users, surface their patterns, and pull the rest of the team toward that practice is running a Stance Two team. The manager who can't is running a Stance One team with a bigger software bill. The metrics below are designed to tell those two apart.


📄 TEAM / MANAGER METRICS

<p>Active usage rate</p> <p>Not licences, not logins. Genuine usage in the flow of work, measured at the action level.</p>	<p>Quality delta on AI-assisted work</p> <p>Error rates, revision cycles, satisfaction on outputs that involved AI vs ones that didn't. Most teams have no idea whether AI work is better, worse, or the same.</p>	<p>Champion leverage</p> <p>Can the manager name the two or three people the rest of the team should be learning from? And is the team set up to learn from how they work, up close?</p>
<p>Rituals integration</p> <p>Does AI appear in standups, reviews, planning, retros? Or is it a private side-tool nobody discusses? It should show up in how the team works, not in a wiki page no one reads.</p>	<p>Knowledge sharing density</p> <p>Are prompts, patterns, and Skills being shared across the team, or is everyone reinventing in isolation? Shared practice is the difference between a team that compounds and a team that doesn't.</p>	<p>Adoption breadth vs depth</p> <p>What percentage of the team uses AI for what percentage of their work? 100% using it for 5% of tasks is a different place from 30% using it for 50%, even if both read as "60% adoption."</p>

✍️ DIAGNOSTIC TEST

If you removed AI tomorrow, would the team's work be visibly worse, or just slightly slower? And the harder version: if the two people most fluent with AI left tomorrow, would the team hold its ground, or would the practice walk out the door with them?

Does AI actually make my job different?



The individual at Stance Two knows whether their practice is being championed or held back, whether the prompt library they built got adopted or filed, whether the next role they're going to do exists in the org's hiring plan or doesn't.

They are watching what the company rewards, and they are making decisions about where to spend the next two years of their career based on what they see.

This is the question most frameworks skip, and it's the most important one given the argument of this paper. If individual roles haven't changed, the company hasn't transformed – it has just installed software. This is also the question your team is asking themselves whether you measure it or not.

It's also where the company's culture call lands personally.

INDIVIDUAL METRICS

<p>Time reallocation</p> <p>What proportion of time is now spent on judgement, edge cases, and synthesis vs the work AI now handles? If the answer is "the same, just faster," the role hasn't changed.</p>	<p>Scope expansion</p> <p>Am I doing things I couldn't do before? Broader ownership, new types of analysis, work that previously required a specialist?</p>	<p>Skill drift</p> <p>What new capabilities have I built? What old ones am I letting atrophy deliberately because AI is better at them than I am?</p>
<p>Confidence calibration</p> <p>Where do I trust AI output, where do I verify, where do I override? Trust should be specific and earned, not blanket.</p>	<p>Future visibility</p> <p>Can I describe what my role becomes, or am I waiting for someone else to tell me? At Stance Three this answer is clear and more interesting than the current job. At Stance One it's a shrug.</p>	

DIAGNOSTIC TEST

Can I describe my job today differently than I would have two years ago, or just say "the same job, with AI"? And the version that matters most: do I believe the company I work for is going to give me the future version of this job, or am I going to have to go somewhere else to get it?

Culture is the ground everything else stands on.

Everything in this paper assumes one thing: that your people can actually move. That the marketing manager can try a new way of working, the engineer can run the latest model and get something wrong, the support lead can build a workflow that fails for a few weeks before it gets good. If they can't, none of the rest matters.

AI rewards one kind of organisation above all others: the one that has shifted from a corporate mindset (process, sign-off, consequence-avoidance) to individuals who are trusted to build solutions that deliver value, whatever their role. That shift is the pre-condition. It looks different in every business, but the question underneath is always the same: when someone tries something and it doesn't work, what happens to them next?

Our Technology paper argued that AI doesn't fix broken processes, it exposes them. One layer down, the same holds: AI doesn't reveal a culture problem, it amplifies one. When the job is being rewritten week by week, fear stops being a background condition and becomes the dominant variable, and fear-based cultures don't produce cautious experimentation, they produce the appearance of it while the work underneath stays exactly the same.

This is the one part of the AI agenda leadership can't delegate. You can hire someone to run the training, buy the tools, commission the strategy. Whether your people are empowered to try, fail, and build is decided at the top and it's decided the first time something goes visibly wrong.

“

When someone tries something and it doesn't work, what happens to them next? That answer is your culture.

You can't fix what you can't see.

If you've tried to buy software in the last eighteen months, you know the feeling: more tools, more "AI-powered" everything, a market louder and less navigable than at any point in software history and a board that wants an AI strategy by Friday.

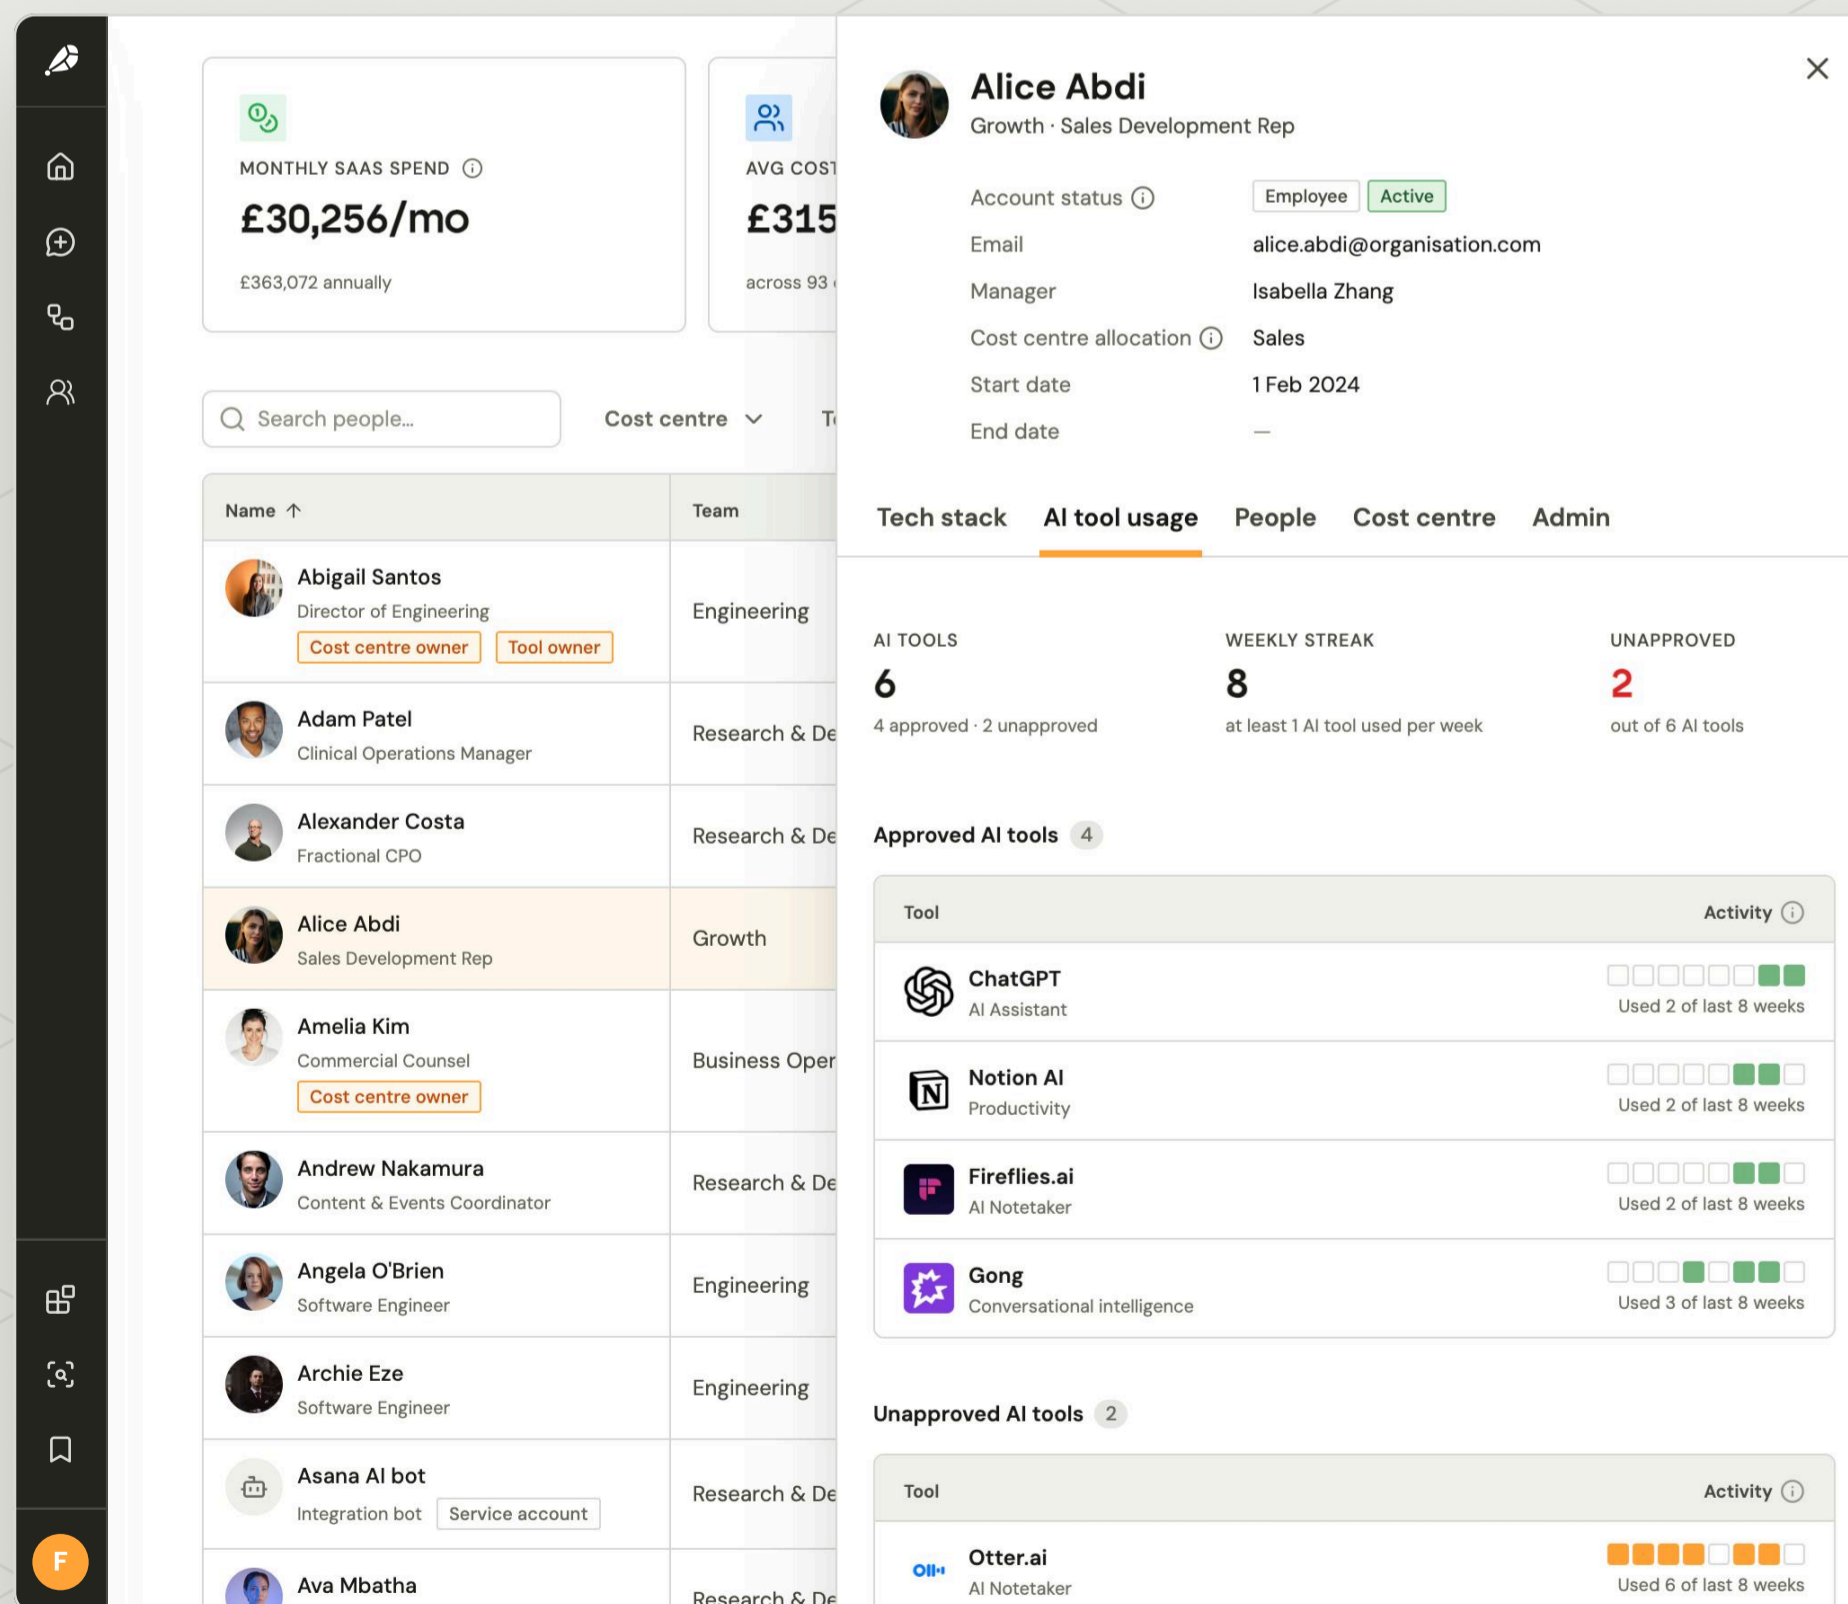
Underneath the noise sit three problems leaders can't get at. The market is impossible to navigate from the inside. The human system AI actually runs on, the identity threat, the incoherent KPIs, the gap between practice and job description, the champions you're backing or quietly losing are invisible from where leaders sit. And the data that would make it visible is scattered across twenty products never built to talk to each other, none of them looking at the people layer.

So the highest-stakes call of the decade gets made on instinct, anecdote, and vendor decks. That's the gap the next pillar of Dragonfly closes, though it's worth being precise about how. Dragonfly won't redesign your roles, run your enablement, or make the culture call for you; those can't be outsourced, and we wouldn't pretend otherwise.

What it does is show you what's actually happening on the ground: where adoption is real, where it's theatre, which teams have changed how they work and which just carry a bigger software bill. The diagnosis is ours. The decision is yours.

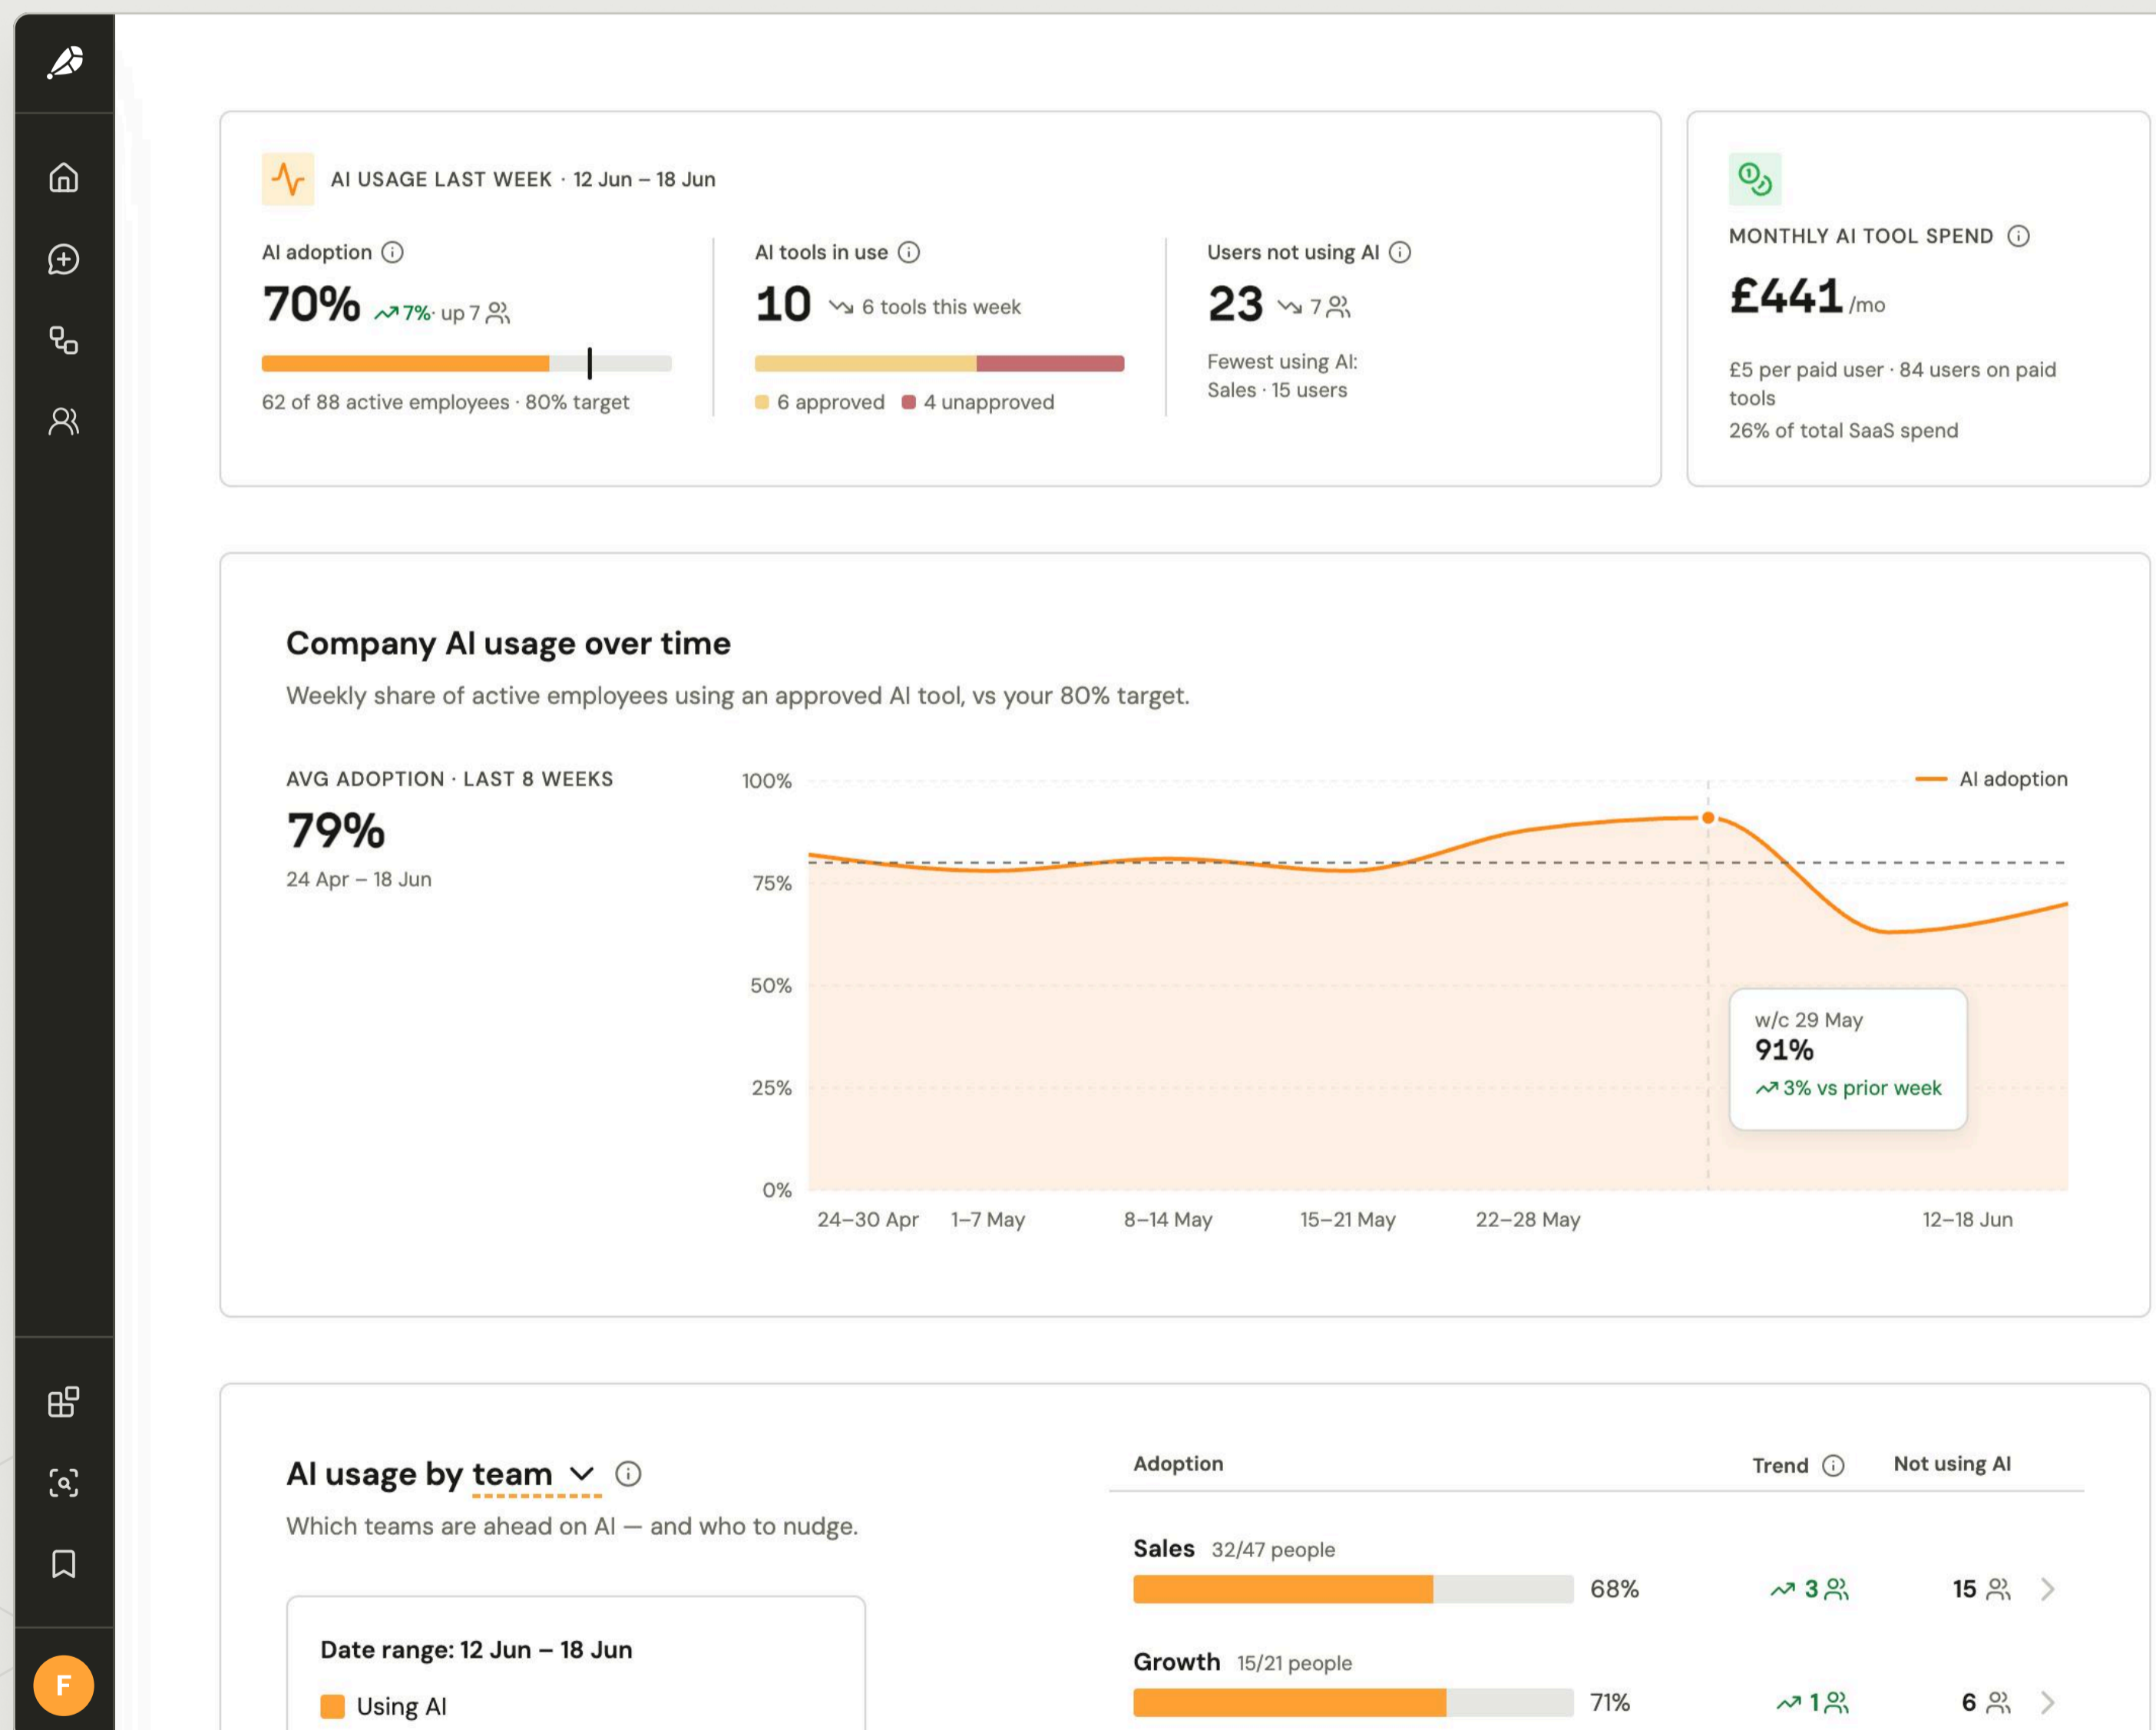
What it does. Connect Dragonfly to the tools your team already uses: your AI platforms, collaboration suites, productivity layer.

It then surfaces the picture no single tool can: who's using what, where practice is concentrated and where it's stranded, which stance each team is actually at. It benchmarks that against comparable companies by size, sector and function, so "are we ahead or behind?" finally has a defensible answer. And it gives every level its own view; output and headcount shift for the C-suite, ritual and knowledge-sharing for managers, scope and skill drift for individuals.



What makes it different? It reads stance and friction, not just usage: where team practice and company policy are pulling apart, where champions are being amplified or held back. It measures impact over activity, built on this paper's diagnostic questions rather than the vanity metrics most dashboards default to.

And it's independent, vendors don't pay to influence what we surface, and the benchmarks aren't weighted toward the loudest voice in the market. The leaders who get this right won't be the ones with the best tools. They'll be the ones who can see what's actually happening, back the champions they already have, and build the org chart that will look obvious in retrospect.



Two pillars down. One to go.

Two papers in, the picture is half-drawn: the stack you're building, and the people who have to make it work. The third pillar 'Strategy' is where those two get tied together: where AI should be deployed, what it should be solving, what guardrails should be in place, and how the business should be making the trade-offs it's now facing every week.


We'll be publishing the Strategy paper next. Together the three pillars of the Dragonfly AI Readiness Framework describe what we believe it actually takes to become an AI-ready business; not a one-off transformation programme, but a continuously evolving operating model that can keep pace with a market that won't slow down.

The next generation of companies won't win by owning more software, and they won't win by training harder either.


They'll win because they've built a human system capable of evolving alongside the tools. A workforce that knows what its job is becoming, a leadership team that can see what's actually happening, and an organisation that has moved through the stances rather than installing software on top of them.

The companies that get to Stance Three won't be the ones with the biggest AI budgets. They'll be the ones that championed the right people, built the practice into how work actually gets done, and rewrote the roles before the market forced them to.

That's what The People Paper is really about. Not training. Not adoption. Not change management, but unlocking the intelligence layer to empower leaders and their team to make the right calls about the work ahead.



[TECHNOLOGY]



[STRATEGY]



[PEOPLE]



Stop guessing, start building.

Map your processes. Score your stack.
See exactly where to invest next, tailored to your
industry, your team, your goals.

askdragonfly.com