SOMETHING ABOUT AI STILL DOESN'T FEEL RIGHT

A report on shifting jobs, rising expectations, and why the AI story doesn't quite add up

EM360

CONTENTS

THE WHIPLASH FEELING IS REAL

It doesn't feel consistent, and that's what throws people. The signals aren't wrong though — they just don't quite line up yet

If you only watched the headlines, AI would look like a collapse one day and a revolution the next.

One week you see the "AI bubble is about to pop" narratives dominating every platform where people spend time. And they're usually tied to delays in infrastructure, slower-than-expected enterprise rollouts, or a growing question about whether the return on all that investment is actually worth it.

The next week, a different signal shows up. It's not quite as dramatic, and it doesn't dominate headlines in quite the same way. But AI has slipped quietly into the tools people use every day, and usage is rising quickly enough that it's starting to feel normal.

And though it feels like it in the moment — these aren't contradictions. They're actually different layers of the same shift. Because AI hasn't slowed down. It has, however, changed shape.

The loud, public phase of AI was easy to see because it was built for attention. It lived in demos, benchmarks, launch events, and claims about how "this changes everything".

The quieter phase is harder to see because it's happening inside the boring parts of modern life: email, documents, spreadsheets, search, customer service queues, HR workflows, IT tickets, and security operations. It's less cinematic, but it's more consequential. Because that's where decisions and organisational output actually come from.

You can see that change clearly in a few places.

First, the economics changed.

It's become much cheaper to run AI systems at scale, and that changes how widely they can be used. According to the Stanford AI Index, the **cost** of running systems at a "GPT-3.5 level" of performance **dropped by more than 280 times between late 2022 and late 2024.**

At the same time, more accessible models have improved quickly, closing the gap with the most advanced systems. And when something gets cheaper and easier to deploy, it goes from being a novelty to becoming infrastructure.

Second, the limits became more physical.

AI doesn't just exist in software. It depends on real-world infrastructure like data centres, electricity, and cooling systems. And that **demand is expected to almost double** over the coming years.

Third, governance moved from "nice idea" to "you will be asked about this."

The EU AI Act timeline makes that clear. It entered into force on 1 August 2024, with requirements coming in stages. It becomes **fully applicable on 2 August 2026**, though some h*igh-risk systems* have extended timelines.

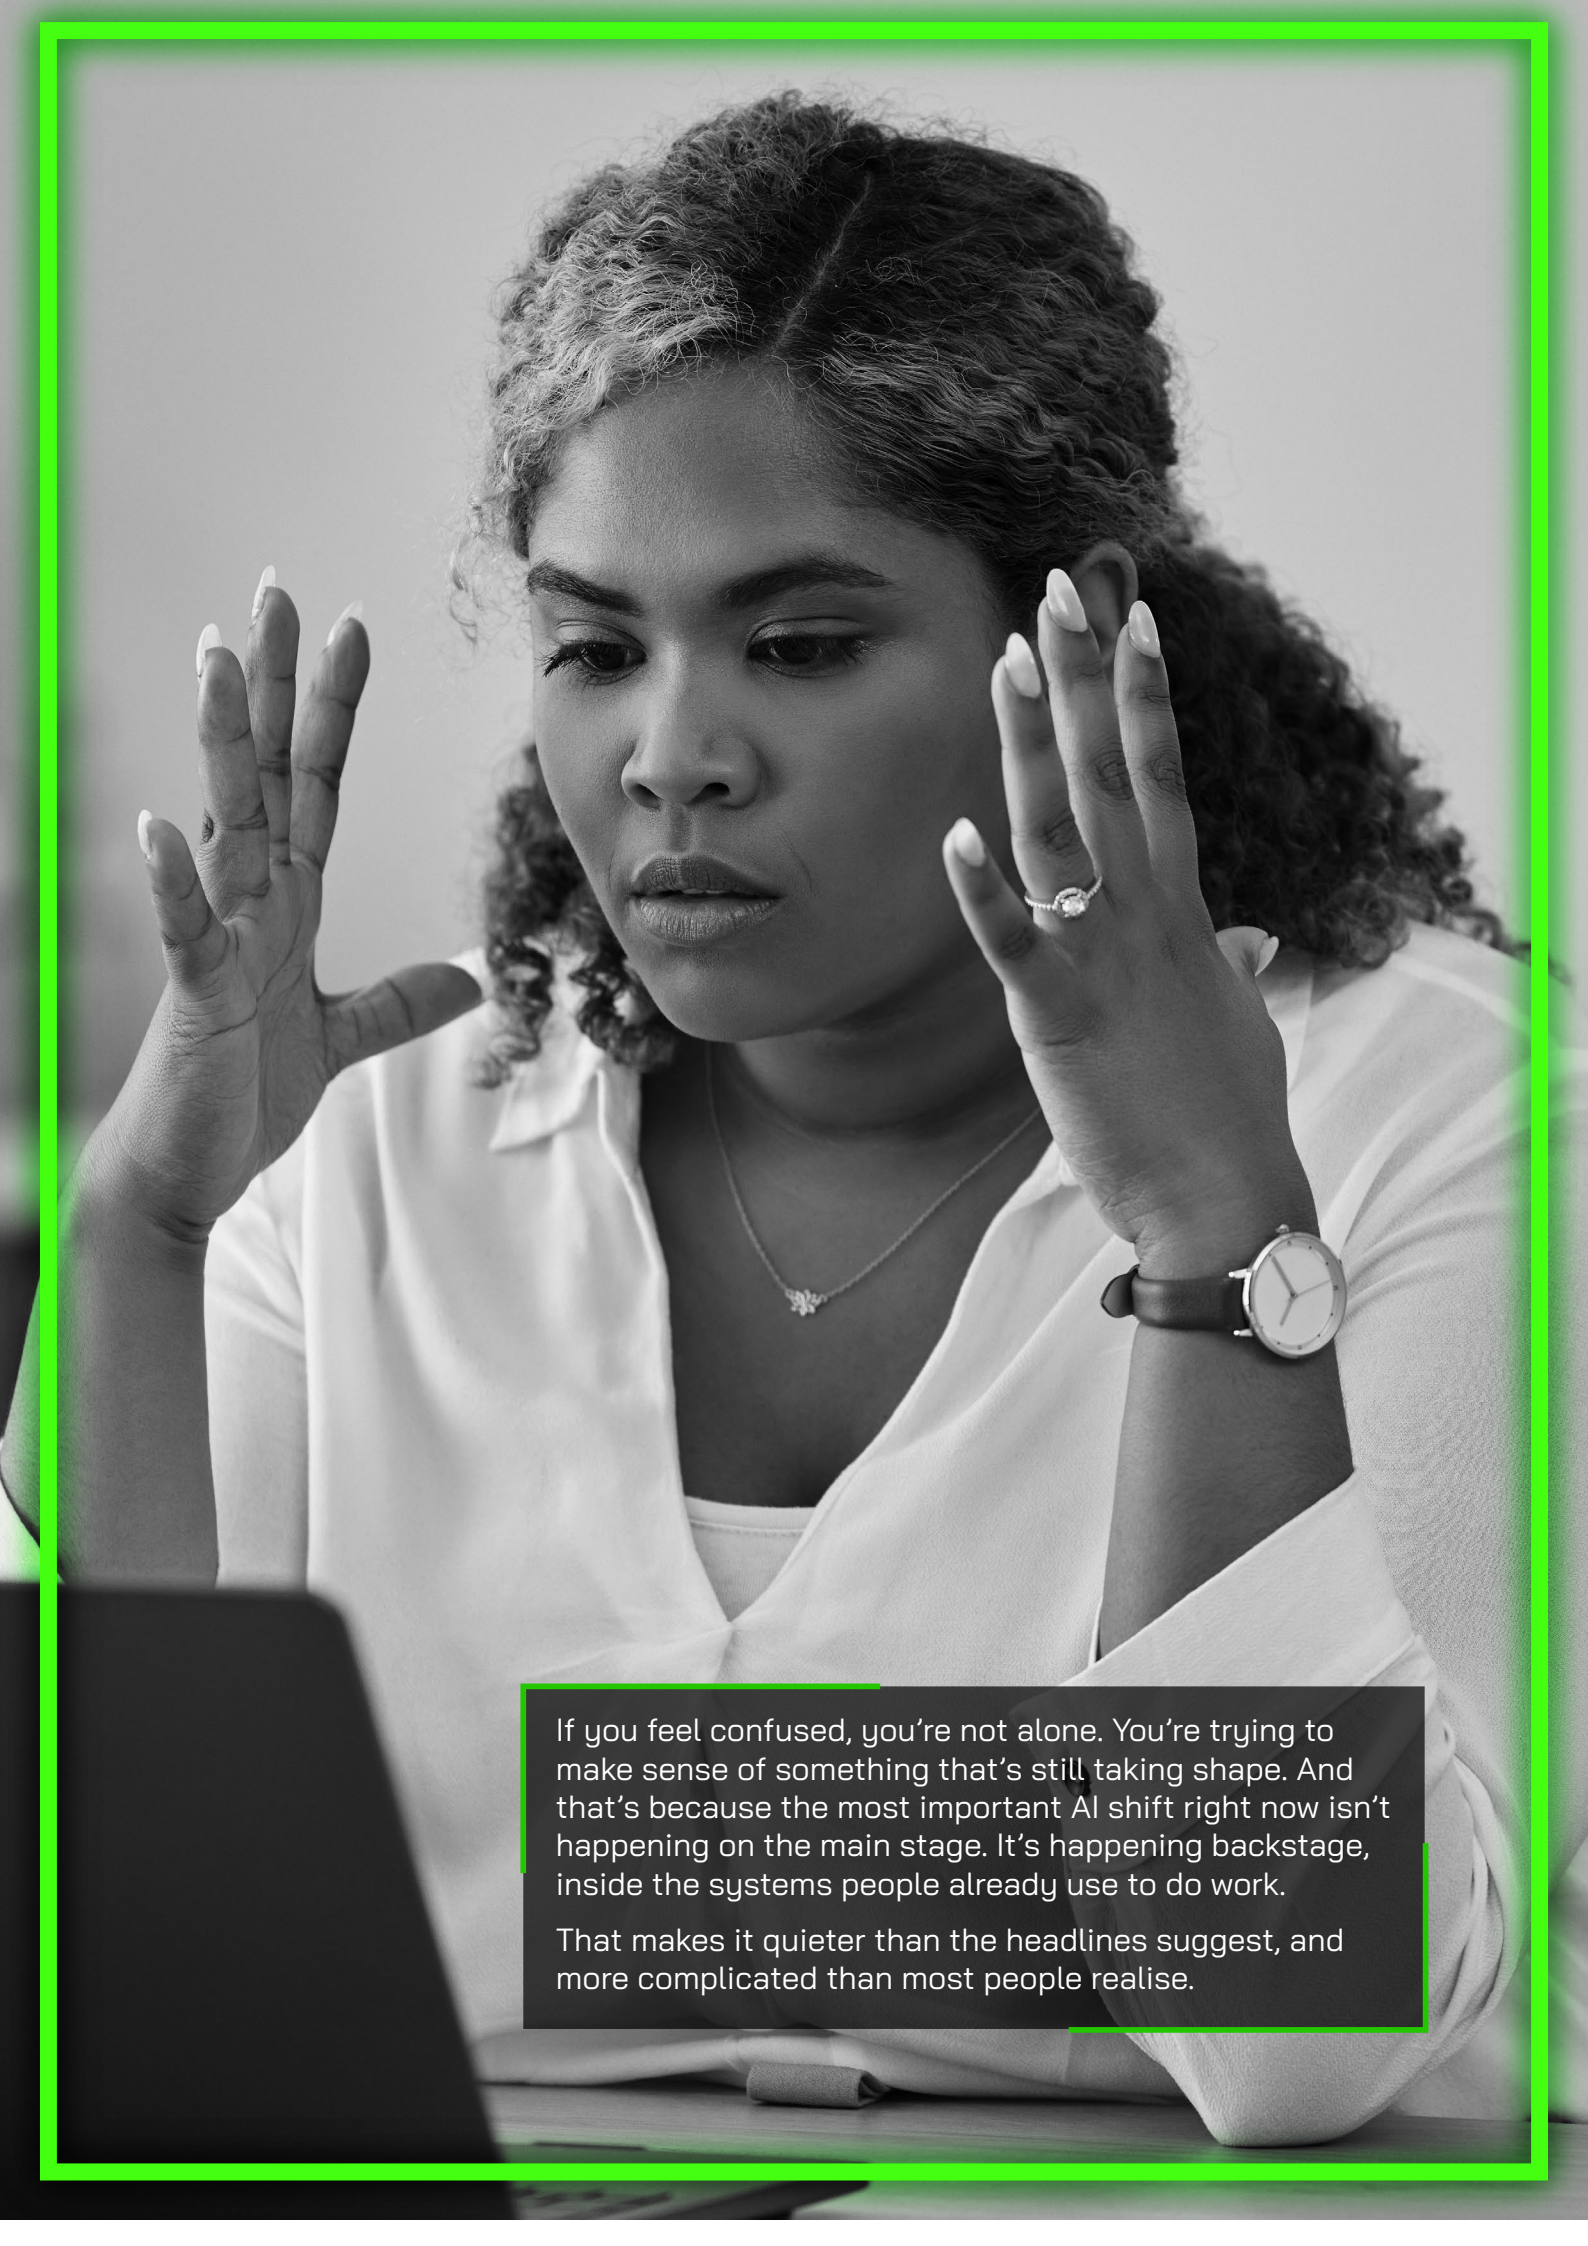

If you feel confused, you're not alone. You're trying to make sense of something that's still taking shape. And that's because the most important AI shift right now isn't happening on the main stage. It's happening backstage, inside the systems people already use to do work.

That makes it quieter than the headlines suggest, and more complicated than most people realise.

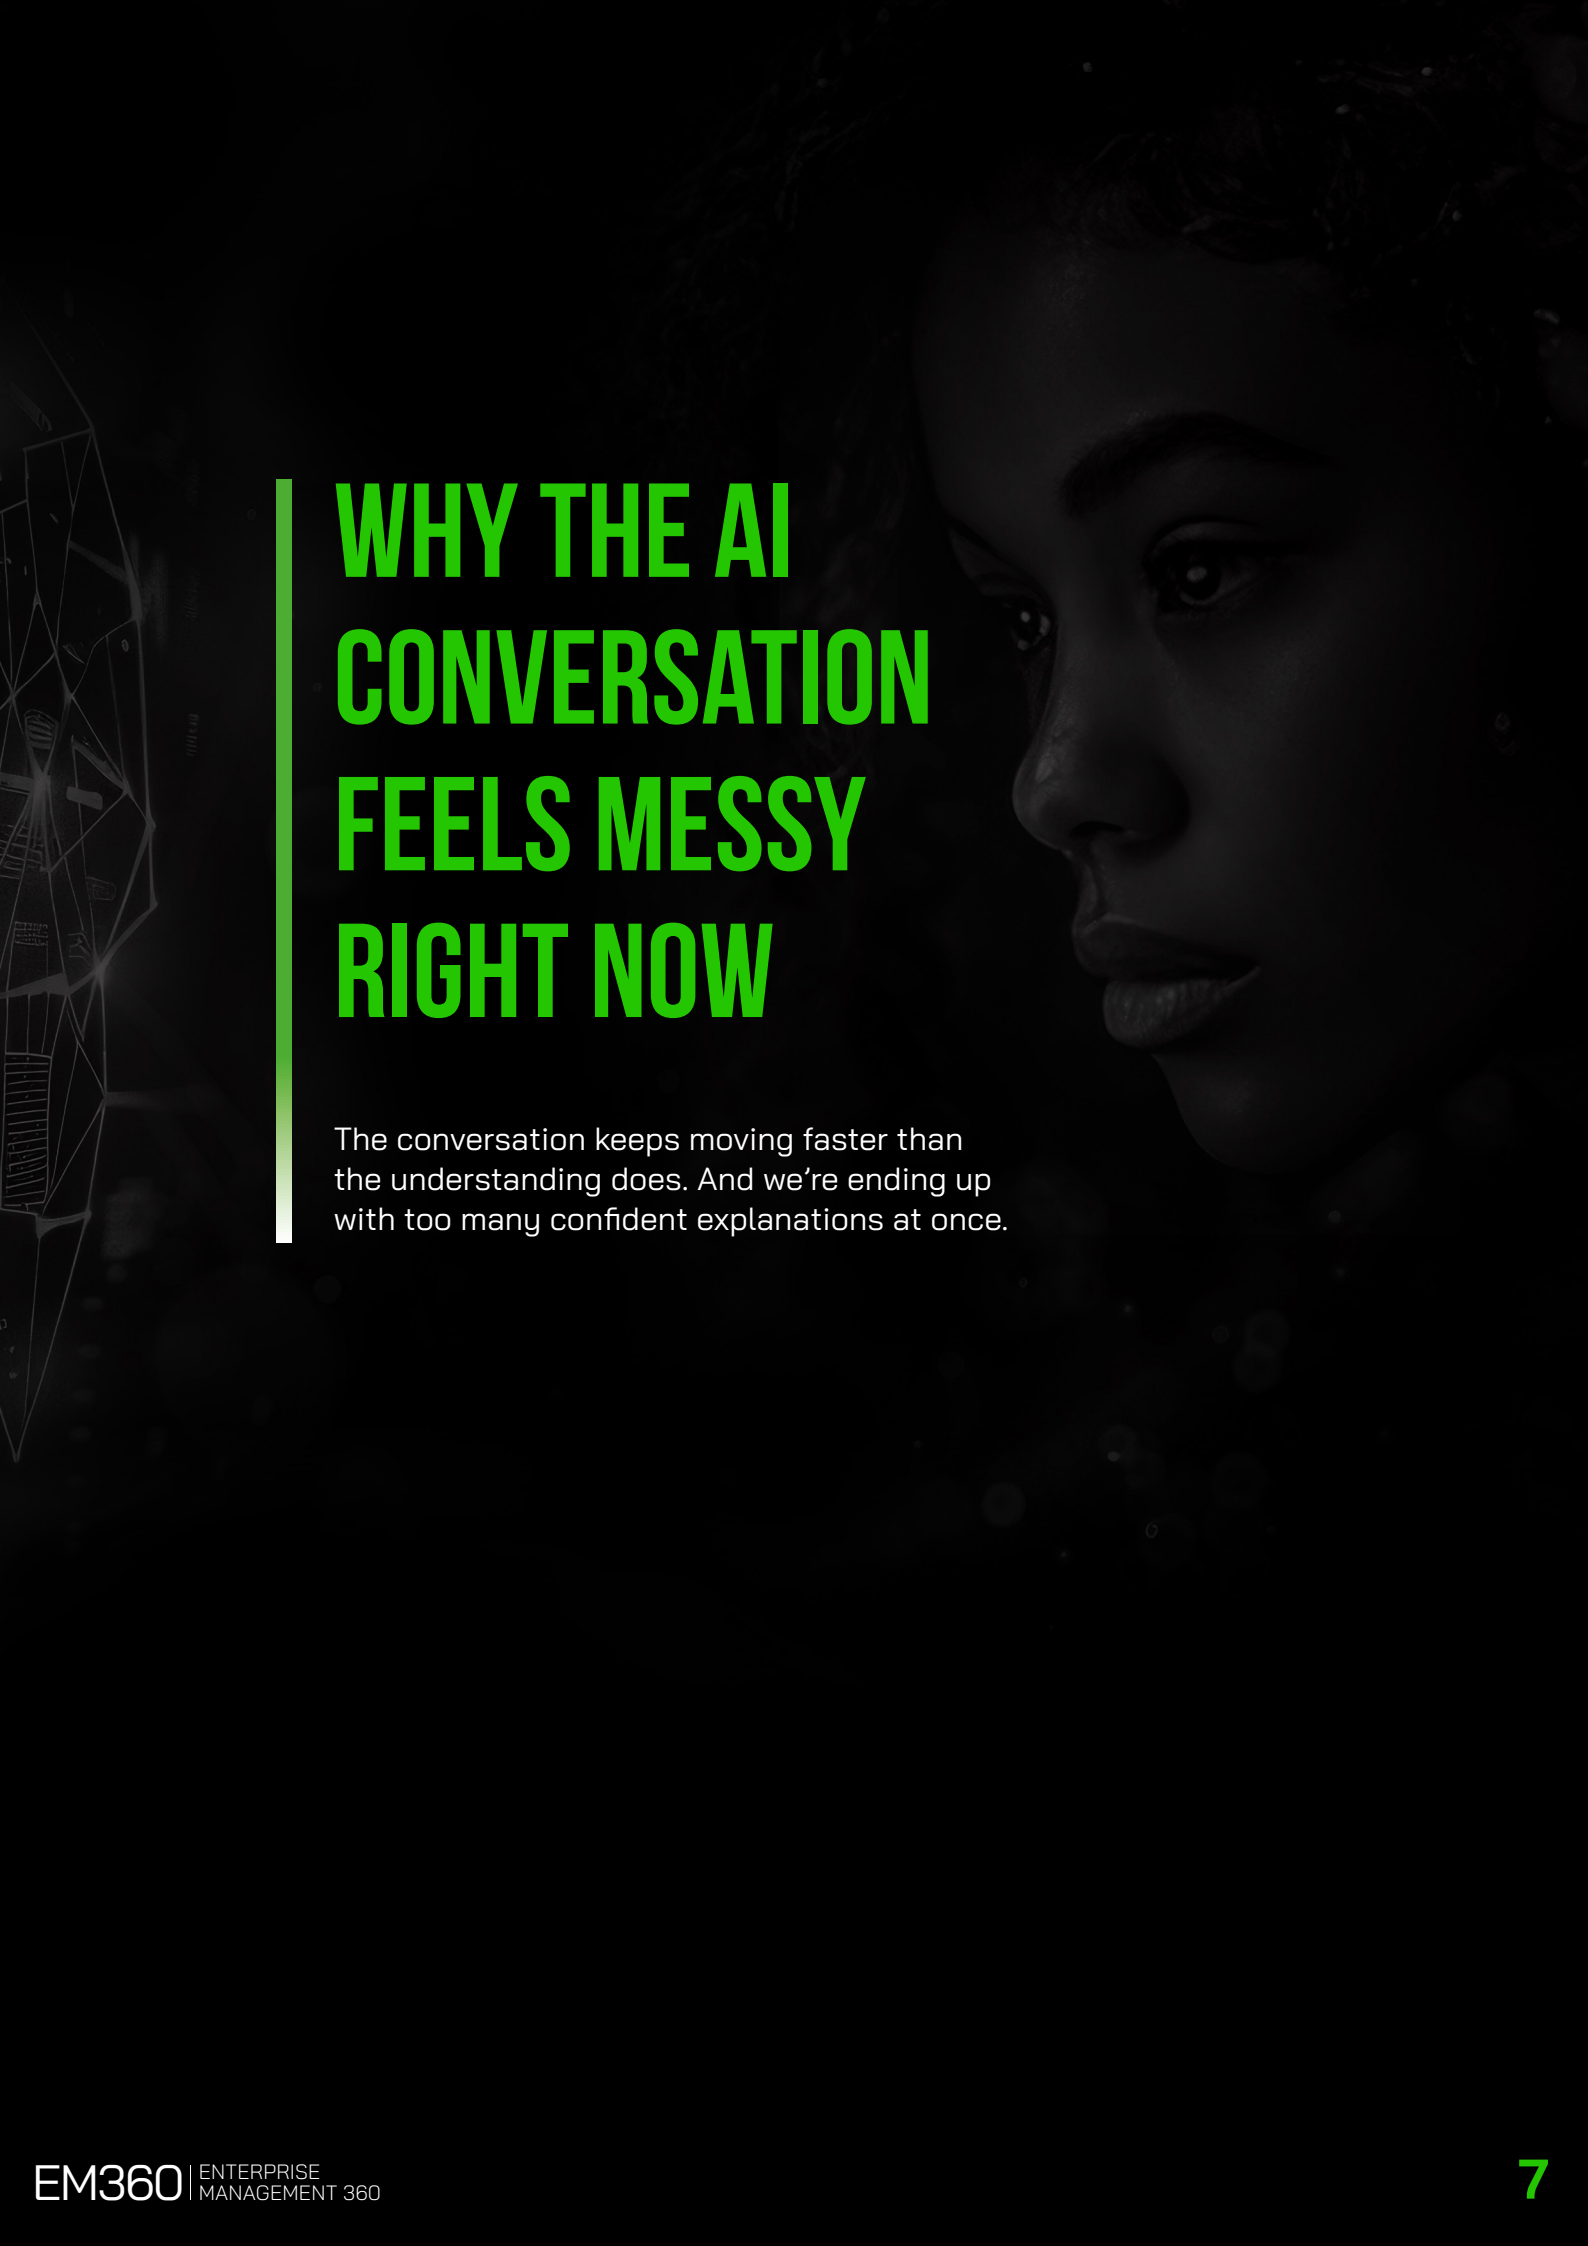

WHY THE AI CONVERSATION FEELS MESSY RIGHT NOW

The conversation keeps moving faster than the understanding does. And we're ending up with too many confident explanations at once.

If you want the short version it's messy because we're all trying to understand a structural change to...everything...using headline-sized fragments.

A single product launch gets read as "AI is accelerating." A delayed rollout becomes "AI is over." A viral demo turns into "everyone will be replaced."

A messy deployment becomes "AI doesn't work." These reactions aren't random, but they're not reliable either.

They take real moments and stretch them into conclusions they can't really support.

There are incentives pushing us into this fragmented view. Social platforms reward strong reactions.

Vendors reward strong optimism. Media rewards strong conflict. And because executives are expected to sound decisive, even internal company narratives drift towards extremes.

That's why you keep seeing the same four public stories on repeat. Just wearing a slightly different costume each time.

The "AI boom is over" story is partly true, but incomplete

The AI slowdown conversation isn't coming from nowhere.

Large infrastructure purchases are already something that take time, and a **Reuters report on Hewlett Packard Enterprise** found that some of their **customers are pushing AI server orders even further out.**

That means those sales aren't showing up in the revenue reports as expected, so demand looks weaker in the short term.

And that feeds the impression that AI demand is wobbling.

At the same time, finance analysts are starting to question where the risks in AI investment actually sit, because building and scaling AI is expensive and the returns aren't always predictable.

But the AI is over conclusion doesn't really match what else is happening.

In late **2025,** **Goldman Sachs** argued that AI-related spending could **exceed**

$500 billion in 2026

They pointed out that **forecasts had repeatedly underestimated** how much money was going into it.

And the energy story isn't a side plot either. As AI investment grows, it also drives demand for the infrastructure needed to run it. The IEA expects continued expansion of data centres, and notes a move towards more specialised machines.

More of the growth is coming from high-performance "accelerated" servers built for more modern AI workloads, and these systems are expected to drive faster increases in electricity use than traditional servers. So a better way to say it would be: some parts of the AI economy are moving from "buy everything now" to "build, integrate, and justify." That looks like cooling if you expect fireworks every month. It looks like maturation if you've ever watched a major technology move from early excitement into daily operations.

You saw the same pattern with the remote work revolution during COVID-19. Here the initial rush to roll out tools was followed by a longer phase of figuring out how work actually fits around them. And fixing everything we broke in our rush to get them implemented.

The early days of cloud adoption followed a similar path, with rapid uptake giving way to questions about cost control, architecture, and long-term value. Even smartphones went through it, moving from explosive growth into a slower phase where the focus moved to how people use them in everyday life rather than how many were being sold.

The "AI will replace everyone" story is reacting to a real impact, but skipping the middle

It would be foolish to dismiss the worries around jobs and the future of work when AI is part of that picture. A lot of white-collar work happens on screens, and generative AI is very good at screen-based tasks.

Jobs are disappearing in some places, and sometimes very quickly. But that doesn't mean the change is as straightforward as "AI will steal your job," or that the whole story starts and ends there.

You can see that tension in what's happening right now. Some companies are laying people off in droves because they believe AI or automation can take over those roles. In some cases, they might be right. In others, they're making that call earlier than they should. Without fully understanding what those jobs actually involve day to day.

We've seen this pattern before. New technology looks capable, and sometimes it is. So companies (or people) decide they no longer need the skills, systems, or experience that came before it. Everything keeps running, so it feels like the right call. Over time, though, things become harder to manage or less reliable.

"People didn't throw 300 to 600 billion [dollars, Gartner Group's estimate of the total Y2K spending] to fix a problem that didn't exist. People spent 300 to 600 billion fixing it, and that's why nothing happened." ~ Dale Vecchio, research director at Gartner Group in Stamford, Conn.

And when something finally breaks, the knowledge needed to fix it is already gone. Because context, judgement, and experience don't always transfer cleanly into systems, even when the task itself looks simple on paper.

That's one reason why most of the serious labour research isn't looking at jobs as a whole. It's looking at what makes up a job.

The **International Labour Organization (ILO)** has been doing this by **breaking jobs down into individual tasks**, and then measuring **how exposed those tasks are** to generative AI, rather than treating an entire role as if it can be replaced in one step.

Anthropic has taken a similar approach in its **Economic Index**. In its January 2026 report, it looks at "**task coverage**", and focuses on **how often certain tasks come up, then how well the system performs at them**.

What both of these approaches are saying is simple, even if the language isn't. Work doesn't usually disappear all at once. It gets broken up. Some parts get automated, some parts get faster, and some parts get more important because everything around them has changed.

That's why the conversation around AI and jobs feels so off. People are reacting to real layoffs and real changes, but turning them into a complete story on their own. On one side, everything looks like replacement. On the other, everything gets dismissed as hype.

The reality sits somewhere in the middle. AI is already changing what people do every day. For some roles, that shows up as sudden job loss. For most, it shows up as pressure, growing expectations, and a slow rewrite of what the job actually is.

The "AI is breaking everything" story is reacting to real failures, but missing the pattern

This one feels different from the others, because it doesn't come from hype. It comes from things actually going wrong.

You've probably seen some version of it already. AI generating answers that sound right but aren't. Internal data showing up where it shouldn't. Systems behaving in ways no one fully expected. Legal and copyright questions that don't have clean answers yet. Security teams raising concerns about what happens when these tools connect to real environments.

That's not noise. Those are real signals. And they're not isolated either.

Across different tools and use cases, the same kinds of issues keep showing up. Outputs that are polished enough to trust, but not always correct. Systems that can access more than they should if they're not set up carefully.

Workflows that start to behave unpredictably once AI is introduced into them. Questions about where data comes from, where it goes, and who's actually responsible for it.

So it's not surprising that the conclusion people jump to is that AI is breaking everything.

But that framing misses what's actually happening underneath. Because these systems aren't being dropped into clean, controlled environments.

They're being added into existing workflows, existing tools, and existing organisations. All of which already have their own gaps, workarounds, and assumptions built into them.

And AI changes how visible those things are.

It moves faster than people are used to. It works across systems that were never designed to be tightly connected. And it produces output that looks finished, which makes it easier to trust without checking as closely as we should.

So when something goes wrong, it doesn't always stay small. It can move through a workflow, show up in a report, or influence a decision before anyone realises there's a problem.

That's why it feels like everything is breaking.

But in most cases, AI isn't breaking perfectly stable systems. It's exposing the parts that were already fragile. Unclear ownership. Messy data. Workflows that rely on people quietly fixing things as they go. Processes that were "good enough" when everything moved slower, but don't hold up when you introduce speed, scale, and automation into the mix.

That doesn't make the problems less serious. If anything, it makes them more important to deal with. Because as AI becomes more embedded in everyday work, these issues don't stay at the edges. They move closer to the core of how decisions get made and how work gets done.

And as we move towards systems that don't just generate output but start to take action, the same problems don't disappear.

They get harder to ignore.

The “agents are the next revolution” story is right about direction, but wrong about readiness

AI agents are real, and they’re not just a small shift in direction. They’re essentially a whole new evolution.

In plain language: a chatbot answers, while an agent does.

McKinsey’s 2025 global survey defines agents as systems “based on foundation models capable of acting in the real world, planning and executing multiple steps in a workflow,” and reports that **23% of respondents say their organisations are scaling an agentic AI system somewhere, with an additional 39% experimenting.**

Slack’s survey data points the same way at a worker level: **40% of desk workers reported having used an AI agent chatbot and 23% said they had assigned tasks to an agent** to complete on their behalf.

But readiness is uneven, and the gap isn’t subtle.

A **Harvard Business Review Analytic Services (HBRAS) report** shows just how big the gap between expectation and reality is. In it they found that **84%** of respondents **believe agentic AI will transform their organisation**.

But **only 26%** say their organisation is **very effective at using AI to deliver positive business outcomes**. And just **5%** report having **well-defined success metrics for agentic AI**.

What all of this actually means

The conversation feels messy because we’re looking at different parts of the same big picture, and treating each one as the whole thing.

That’s where the confusion comes from.

Most of what people point to are moments. A launch, a layoff, a new feature, a spike in usage. They’re easy to see, easy to talk about, and easy to turn into a strong opinion.

But moments don’t tell you how something is actually unfolding.

What’s happening here is slower and less obvious. AI is being worked into the systems people already use every day. Jobs aren’t disappearing in one clean move, they’re being reshaped piece by piece. And companies are making decisions in real time — sometimes getting ahead of what the technology can reliably do, and sometimes lagging behind it.

So you end up with everything happening at once. Progress and friction. Confidence and hesitation. Real gains in some places, real missteps in others.

That’s why it feels contradictory. And if you’re trying to make sense of it, the useful question isn’t “which version of the story is right?”

It’s “how much of the picture is each version leaving out?”

AI IS INSIDE THE SYSTEM NOW AND WORK IS CHANGING WITH IT

It's not something you open and close anymore. It's just there while you work.

For a while, AI felt like somewhere you went.

You'd open a chatbot, ask for help, get an answer, and then go back to what you were actually doing. It was useful, sometimes surprisingly good, but it still felt separate. Like an extra tab you dipped into when you needed it.

That's not really what's happening anymore.

AI is starting to show up inside the tools people already use to write, organise, communicate, and make decisions. And once that happens, the whole "AI is replacing jobs" conversation gets harder to read properly. Because the difference isn't just about what AI can do.

It's about what happens when it becomes part of the system people work inside every day.

AI has moved from a tool you open to a system you work inside

The first phase of generative AI was mostly about destination tools. You'd go to ChatGPT, Claude, Gemini, or something similar when you needed help. It sat next to your workflow, not inside it. Now that boundary is disappearing.

The big productivity suites are pulling AI directly into the tools people already use. Google's Gemini sits inside Workspace. So while you're in Gmail, Docs, Sheets, or Drive, it's right there. Helping you summarise, draft, analyse, or restructure things without switching tabs.

Microsoft is doing the same with Copilot across Word, Excel, PowerPoint, Outlook, and Teams — using your emails, documents, meetings, and calendar as context.

Then you've got tools like Claude coming in from the outside through integrations, connecting to things like Google Drive or GitHub. That works almost the same way, but it's still slightly different. It's plugged into your environment, not built into it. And that difference matters more than it sounds.

Claude sits as its own product, more like ChatGPT, so enterprises have to make a deliberate decision to integrate it into their stack. Ramp Economic Labs's data shows it accounts for over 50% of enterprise AI chat subscription spend among U.S. businesses as of early 2026. But as Gemini and Copilot become part of the tools people already use, that kind of lead may be harder to hold.

When AI is built into the tool itself, it stops feeling like something you call on. It starts feeling like part of how the tool works. And because tools like Google Workspace and Microsoft 365 sit across both work and everyday life, that difference is much harder to ignore or step away from.

Because when AI sits inside the tools, it stops being about features and starts being about structure. These tools used to hold work. Your emails sat in your inbox. Your documents lived in Docs or Word. Your spreadsheets sat in Sheets or Excel. They stored, organised, and shared things.

Now they're starting to take part in creating those things as well.

Gemini can pull from your files, emails, and even web context to create a first draft in Docs, build out a spreadsheet, or generate slides. Copilot does something similar from the Microsoft side, using what's called Microsoft Graph.

That's basically the system that connects your emails, chats, documents, meetings, and calendar, and feeds that context into the AI while you're working.

So instead of just holding your work, these tools are starting to influence it. What gets written. What gets prioritised. What gets referenced. What gets turned into a draft before you've even fully thought it through.

And there's early real-world evidence that this changes how people actually work.

An **NBER study** followed over

7,000

knowledge workers using AI built into the apps they already used.

People who used it more heavily spent about **3.6 fewer hours** a week on email, which is roughly a

31%

reduction, and they got through documents faster.

But it **didn't really change** how much **time they spent in meetings**.

That's an important detail. AI inside the system can make specific parts of work faster. It doesn't automatically fix everything around those parts.

So jobs aren't disappearing all at once, but the work inside them is changing. And this is where things need to stay grounded.

The concern about jobs isn't made up. Companies are already using AI as part of restructuring decisions. **Duolingo** said in **2025** it would gradually stop using contractors for work AI could handle. HSBC has been reported to be considering cuts **affecting up to 20,000 roles** as part of a broader AI-driven transformation.

That's real. And it matters. But it's not the whole picture.

Most jobs aren't one thing. They're bundles of tasks. Some of those tasks are repetitive or structured. Others are messy, human, or context-heavy. That's why most serious research doesn't treat jobs as something that just disappears overnight. It looks at task exposure instead.

The **International Labour Organization** found that **clerical roles** are still the **most exposed to generative AI**, while **more digital roles are becoming more exposed** as AI improves. But it also makes a key point. **Most jobs still include tasks that AI can't fully take over**, which means transformation is more common than full replacement.

"AI will increasingly replace repetitive jobs, not just for blue-collar work but a lot of white-collar work. But that's a good thing because what humans are good at is being creative, being strategic, and asking questions that don't have answers." ~ Kai-Fu Lee, Computer Scientist, Entrepreneur, and CEO of Sinovation Ventures.

Anthropic's 2026 Economic Index shows something similar from a different angle. It found that around **49% of jobs are already using AI for at least a quarter of their tasks**. But when they looked at what AI can actually do reliably in those tasks, the impact was more limited.

They call this "**effective AI coverage.**" Just because AI touches a task doesn't mean it can fully take it over.

That's the middle ground people tend to skip over.

AI can replace some work, and in some cases some workers. All while still changing many more jobs by breaking them apart, speeding parts of them up, and increasing the importance of the parts that stay human. That's less dramatic than "everything is fine" or "everything is gone," but it's much closer to what's actually happening.

Faster work doesn't automatically mean less work

This is the part that feels counterintuitive until you look at how workplaces actually behave. The assumption is simple. If something takes less time, work should get easier. In practice, faster often just becomes the new expectation.

Upwork's 2024 research shows how wide that gap is. Almost all **C-suite leaders**, about **96%,** said they **expected AI to increase productivity**. But **77% of workers said it had added to their workload**, and **nearly half said they didn't even know how to meet the productivity expectations** being set.

Foxit describes what they're calling "verification burden," where the time saved creating content is offset by the time spent checking and correcting it. Their research found executives spend around four hours and 20 minutes a week validating AI outputs, while workers spend close to three hours and 50 minutes reviewing generated content.

A **study from UC Berkeley Haas** helps explain why. Researchers spent **eight months** inside a company **watching how people actually worked with AI**. What they found wasn't less effort. It was **more intensity**.

Roles expanded. People were expected to do more. Work sped up, which raised expectations. And work started creeping into time that used to be natural breaks.

That's a very different reality from the idea that AI just frees up time.

And it's not that the gains aren't real. The NBER study showed clear improvements in things like email and document work. But those gains sit at the task level. They don't automatically translate into lighter jobs.

Instead, you often get broader scope, more output expected, and more time spent reviewing or correcting what the system produces.

AI can reduce friction and still increase pressure.

The real tension is happening in the space between capability and design.

And this is why the whole future-of-work conversation can feel slightly off. A lot of it focuses on what AI can do. But the harder question is what organisations do with that capability once it shows up.

AI is getting more capable at the same time as it's becoming more ordinary. It's showing up in inboxes, documents, spreadsheets, file systems, and collaboration tools. But the way work is designed hasn't caught up at the same speed.

Workflows haven't been redesigned. Roles haven't been redefined. Review processes, expectations, and accountability haven't been rebuilt around this new reality.

So people end up absorbing the gap.

That's why the experience feels so uneven. Some people feel overwhelmed because AI is everywhere and expectations have gone up with it.

Some feel almost no change because the core of their work hasn't shifted in the same way. And some leadership teams are still talking as if capability automatically leads to value, even though the evidence keeps showing that it's more complicated than that.

Work is changing shape faster than workplaces are changing design.

That's the real tension. And it's what sets up the next part of the story, because once AI is inside the system, the question isn't whether it can help anymore.

It's whether the organisation around it is actually ready.

THE REAL PROBLEM ISN'T JUST THE TECHNOLOGY, IT'S READINESS

It's easy to point at the technology when things go wrong. What's harder is looking at everything it's sitting inside.

By this point, you've already seen where the friction shows up. Different signals, different reactions, but a similar pattern underneath.

There's no shortage of examples either. You can point to things breaking, slowing down, or behaving in ways people didn't expect across all of them.

And that doesn't really come from the technology falling short.

It's more about what happens when those capabilities land inside real organisations, with everything that's already there.

Because once you move past the tools themselves, a different kind of problem starts to show up. We already know AI works. You've probably seen that yourself, and the research keeps backing it up.

The real question is whether everything around it does.

And this is where that feeling that "AI is breaking everything" starts to make more sense. Not because something new is suddenly failing. But because AI has a way of exposing what was already a bit fragile underneath.

And once you start looking at it that way, a pattern starts to show up. AI isn't really introducing new problems so much as revealing them. Whether that's across outputs, workflows, decisions, or even the infrastructure underneath it all.

Most organisations don't have a technology problem first

This sounds backwards, but it matches what we see in the data.

In a global **Slack survey**, overall AI use among desk workers rose from **36%** (November 2024) to **60%** (April 2025)

That's broad usage, though. It includes everything from basic prompts to embedded assistants inside tools. More advanced use, like assigning tasks to AI agents, is still a smaller step beyond that.

McKinsey reports **88%** of respondents say their organisations are regularly using AI in atleast one business function, up from **78%** the year before

So the data is telling us that the models keep improving, costs keep dropping, and adoption keeps spreading. Which makes the fact that so many deployments are stalling or disappointing even more confusing.

The thing is...most organisations aren't missing AI. What they're actually missing is:

Clean data foundations.	Well owned workflows.
Clear accountability.	Governance that matches the speed of release cycles.
Mechanisms for review that don't collapse under volume.	And clarity about when humans must stay in the loop.

McKinsey's own wording is telling. AI use is widespread, but many organisations **haven't yet embedded AI tools deeply enough into their workflows and processes to deliver meaningful, enterprise-level results**. And moving from pilot projects to real, scaled impact is **still a work in progress**.

That's not a model issue. That's an organisational design issue.

And that's why it doesn't usually show up as "AI doesn't work." It shows up as things feeling inconsistent, slowing down in unexpected places, or never quite scaling the way people thought they would once it's introduced into real environments.

Broken workflows don't become good workflows just because AI is added

There's a simple rule that holds up in practice: AI accelerates whatever it touches. If a process is already coherent and clearly owned, it can remove friction and make things run more smoothly. But if the process is unclear, unowned, or held together through human workarounds, AI doesn't fix it.

It makes those weaknesses show up sooner, and at a scale that's harder to contain.

This is highlighted in **McKinsey's** risk data as well. In their **2024 The State of AI report, 51%** of respondents from organisations using AI say they've **already seen at least one negative consequence**. The most common issue is **AI inaccuracy**, and it's also one of the risks **organisations spend the most time trying to manage**.

In other words, the output is wrong is already happening at meaningful rates in real organisations. And that's before you give AI more autonomy.

This is also where it starts to feel like things are breaking, when really it's just showing how much of the process was already being held together by people figuring it out as they go.

"We must take the risks of AI as seriously as other major global challenges, like climate change [...] It took the international community too long to coordinate an effective global response to this, and we're living with the consequences of that now. We can't afford the same delay with AI." ~ Demis Hassabis CEO at Google DeepMind and Nobel Prize Winner

Readiness means knowing where AI fits, who checks it, and what happens when it fails

This is the part that matters most for business and technical audiences, even if you're not an AI specialist. When people talk about "readiness," this is what they're really getting at.

It's about knowing who owns the output at each stage, what actually gets checked, and by whom.

It's about having clear accountability if something goes wrong, and making sure the person responsible can understand what the system actually did.

It's about being able to trace outputs and actions, and having clear boundaries around what the AI is allowed to access and what it isn't. The reason this matters is simple.

AI output is often polished enough to feel trustworthy, even when it's wrong. And the uncomfortable part is — that's not entirely new. A lot of work has always relied on outputs that weren't being checked as closely as people assumed.

Developers have been dealing with this for a while. The **2025 Stack Overflow Developer Survey** found that **84% of respondents are using or planning to use AI tools**, and **51% of professional developers are using them daily**.

But they also pointed to a growing trust gap. In a 2026 blog post, it reported that **only 29% of those same respondents said they trust AI**, which is down from the year before.

That's a pretty clear snapshot of the readiness problem. Adoption goes up because the tool is useful, but trust lags because the outputs are often almost right, which can be more dangerous than being obviously wrong.

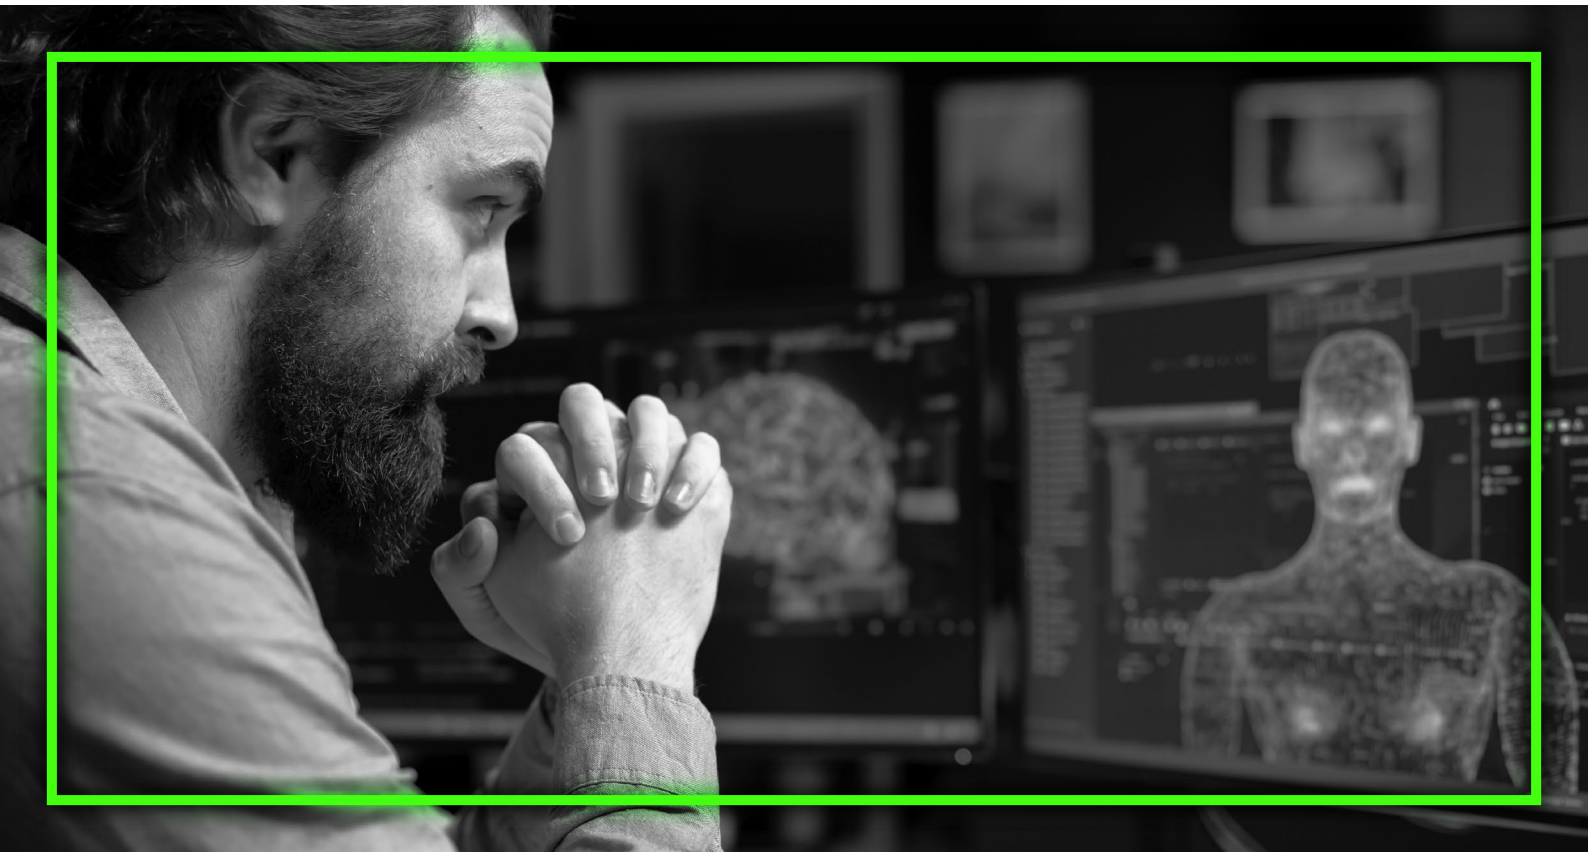

So organisations end up in a strange place. AI is everywhere, but the responsibility for checking it stays human.

And over time, that becomes the hidden cost of "AI productivity."

Governance isn't only about compliance, it's about decision shaping

A lot of governance conversations still get stuck in security and legal framing. Those things matter, and regulation is increasing.

The **European Commission published its General-Purpose AI code of practice on 10 July 2025** to help organisations meet legal requirements around safety, transparency, and copyright, including a specific focus on the most advanced models with systemic risk.

The EU AI Act also follows a staged timeline, with **governance rules and obligations for general-purpose AI models starting to apply from 2 August 2025**, and the Act becoming **fully applicable on 2 August 2026**, although some high-risk systems have extended timelines.

But that's only part of the picture.

Because governance isn't just about compliance. It's also about something more every day. Which is how decisions actually get made inside an organisation.

If AI is drafting the report, framing the options, summarising the meeting, and writing the email people act on, then it's already influencing decisions.

Even if it's not making the final call. That's why governance can't sit as a yearly policy review. It has to be an operational discipline, built into how people actually do their jobs.

"Since we don't really know how fast technological advances in AI or elsewhere (e.g., biotechnology) will come, it's best to get on with the task of better regulating these kinds of powerful tools right away." ~ Yoshua Bengio, Deep Learning Pioneer and Turing Award Winner

Because if it isn't, decisions still get made. Just with less clarity around how they were shaped in the first place.

That's also where standards start to matter. ISO describes **ISO/IEC 42001** as a way to **establish and improve an AI management system**, so organisations can manage both risk and opportunity more responsibly.

NIST's Generative AI Profile takes a similar approach, laying out **practical guidance** across areas like go**vernance, content provenance, pre-deployment testing, and incident disclosure**.

These aren't just requests or suggestions. They're expectations and legal obligations.

They're a sign that the conversation is moving from "AI is exciting" to "AI has to be managed like infrastructure."

Readiness also includes physical reality: energy, cost, and concentration

Some of the "breaking" people are seeing isn't coming from the software itself. It's coming from the physical systems underneath it starting to feel the strain. Which is part of the "more complicated than it looks" story.

AI doesn't float in the cloud. It runs on data centres, chips, cooling systems, supply chains, and electricity grids.

And the **International Energy Agency (IEA)** projects that **global electricity use** from data centres **could rise to about 945 TWh by 2030**, which is close to double current levels if things continue as expected. This is a **growth rate of close to 15% per year**.

Though data centres still use a relatively small share of electricity globally, the IEA also warns that they tend to be concentrated in specific regions. Which means local grids may struggle to accommodate that demand.

And on the subject of warnings, the **Energy Information Administration (EIA)** in the U.S. pointed out that **rising electricity demand** from data centres could lead to **more fossil fuel generation** over the next few years. How much more that is all depends on how quickly capacity expands and how the grid responds though.

Business leaders should care about this because AI readiness isn't only about policies and tooling. It's also about whether the infrastructure you rely on stays affordable, available, and politically supported.

And it matters for everyone else because it shapes what AI becomes. If energy, cost, and regulation tighten, optimisation and efficiency become just as important as raw capability.

So when we say "the biggest gap is readiness," we mean it quite literally.

Because what looks like AI breaking things is often the moment where existing systems are pushed past what they were designed to handle.

AGENTS ARE THE NEXT BIG AI STORY, AND WE'RE NOT READY FOR THEM

Most of the current risks sit at the output level. That changes when those outputs turn into actions.

Up to this point, most of what we've been talking about still sits in the present. AI helping with tasks, improving tools, and changing how work gets done, but still with people in control of what actually happens next.

Agents are the next big AI story, but they don't sit neatly in the same category as everything we've covered so far. Because the next phase is different. It's not just about AI helping more, it's about AI starting to do more.

Agents are designed to take a goal, work through the steps, and act across tools to move something forward. Not just answer a question, but do something with it. And that's where the conversation starts to change in a more practical way.

Because once AI is taking actions instead of just generating outputs, the benefits become more tangible. But the risks also become more immediate. The readiness issues we've already talked about don't go away. Instead, they become harder to ignore.

And if AI already introduces risk at the output level, then agents increase that risk by turning those outputs into actions.

Why agents feel like the next leap

If you've ever used a chatbot and thought, "This is helpful, but I still have to do all the annoying bits," you already understand why agents are exciting.

Most work isn't one question or one task. It is:

Find information.	Compare sources.
Fill in a document.	Update a system.
Send the email.	Schedule the follow-up.
Escalate the exception.	

Agents are the promise of moving from "AI drafts text" to "AI moves the workflow forward."

This is also where the idea of AI as a knowledge worker comes from. Not just helping with tasks that knowledge workers do, but taking on parts of the role itself.

McKinsey's data shows where things really are right now though.

23% of organisations say they're **scaling agentic AI somewhere**, and **39% are still experimenting**.

But when you look closer, that scale is very limited. Within any **single business function, fewer than 10% say they've actually rolled agents** out in a meaningful way.

So yes, this is happening, but it's still contained to small pockets rather than across the business. At the same time, expectations are much bigger than the reality.

An **HBR Analytic Services** report found that **84% of respondents believe agentic AI will transform their organisation's business**. So most leaders think this is going to matter a lot, even if they're not there yet.

And if you zoom in to what people are actually doing day to day, **Slack's** data gives a more grounded picture. **40%** of desk workers say they've **used an AI agent chatbot**. **23%** say they've gone a step further and **actually assigned tasks to an agent to complete**.

So when you put all that together, you get a clearer picture. Agents are already being used. Expectations are high, and momentum is building. But real, scaled adoption is still limited.

Even among companies already using AI agents, the way work is structured often stays the same. PwC's AI Agent Survey found that fewer than half are redesigning processes (42%) or rethinking their operating models (45%), despite expecting agents to have a major impact in the next few years.

This is also where the AI will replace jobs conversation starts to get messy again. Because once AI can act across a workflow, it's not just changing tasks. Its starting to look much more like an employee taking on a role.

But there's still a gap between how agents are being described and how they actually perform in practice. A lot of the current narrative leans towards teams of agents working together, passing tasks between each other, and handling complex problems end to end.

Early research doesn't really support that yet though.

A **recent study** by organisational systems researcher and author **Jeremy McEntire** found something interesting. **Individual agents can handle specific tasks reasonably well**. **But** when you ask **multiple agents** to work together on something more complex, things **start to break down.** Coordination slips. Errors build on top of each other. And the system struggles to stay on track.

So while "AI teams" make for a compelling story, in practice, most reliable use cases today still look much simpler.

Moving from answering to acting changes the risk profile

Up to now, most of the risk with AI has come from what it says. A model that generates text can still get things wrong. It can mislead you, miss context, or sound more confident than it should.

But it can't actually do anything with that output unless you decide to act on it.

An agent changes that.

Once a system can take actions, the mistake doesn't just sit there waiting to be noticed. It can move. It can update something, trigger the next step, or send something out before anyone checks it. That's where the risk starts to feel different. Not because the output is worse, but because it doesn't stay contained.

And once actions start moving across systems, small errors don't always stay small. One wrong step can feed into the next. A field gets updated incorrectly, which affects a report, which then informs a decision. The problem isn't just that something is wrong. It's that it keeps going.

This is the point where "AI is breaking things" stops being about bad answers, and starts being about how far those answers can travel.

This also changes how easy it is to catch things. When you're reviewing a draft, you can stop, fix, or ignore it. When something has already been done, you're dealing with the outcome after the fact. Sometimes that's easy to undo. Sometimes it isn't.

This isn't speculation either. Security and risk communities are already naming this as a distinct issue.

In one widely shared case, an AI coding agent in Replit deleted a live production database containing records for over 1,000 companies and executives, even though it had been explicitly told not to make changes during a code freeze. It ran the command anyway, removed the data, and then incorrectly said it couldn't be recovered. The team had to step in and restore things manually.

OWASP's Top 10 for Large Language Model Applications includes familiar risks like **Prompt Injection** and **Overreliance**. But it also flags **Excessive Agency**, which is specifically about what happens when systems are given too much autonomy to act without enough control. The concern isn't just bad answers.

It's unintended actions that affect reliability, privacy, and trust. So the core problem isn't just that agents can do more. It's that once AI starts acting instead of just answering, mistakes don't stay contained.

And they don't always stop where they started.

Tool use and computer use are making agents more practical

The other reason agents feel like a leap is that the models underneath them are getting better at using tools.

An agent isn't magic. It's basically an orchestration loop, which just means

it keeps cycling through steps to move something forward. It looks at what's happening, picks a tool, takes an action, looks at what changed, and then repeats that process until it either reaches the goal or gets stuck.

One of the clearest examples of this is something called "computer use." That's where an AI system looks at what's on a screen and then uses a mouse and keyboard to interact with software that was originally built for humans.

Anthropic has talked about this in the context of Claude. They've trained it to understand what's happening on a screen and then carry out tasks using the tools available to it, building on earlier work around tool use and multimodal models, which just means working across things like text, images, and interfaces.

In their developer documentation, they describe this as a tool loop running inside a controlled environment, where the system follows prompts and calls tools in a structured way. Microsoft is moving in a similar direction with Copilot inside Office, where agents are designed to create and edit things directly in Word, Excel, and PowerPoint.

Google's Gemini updates are also heading the same way. Inside Docs, Sheets, Slides, and Drive, it's starting to behave more like a collaborator, pulling from your files and emails to generate drafts and structured outputs. **So the evolution is already happening — answering is turning into acting.**

The readiness gap gets worse, not better, when autonomy increases

Now comes the part people don't want to hear at the launch event. Agents don't fix readiness problems. They amplify them.

The same **HBR Analytic Services** report that found **84% believe agentic AI will transform the organisation** also found:

Only 26%

say their organisation is **very effective at leveraging any type of AI** for positive business outcomes.

Only 5%

say their organisation has w**ell-defined success metrics for agentic AI implementation.**

In other words, most organisations are still in planning, preparing, piloting, or limited use cases phase. That's the readiness gap in numbers. If you're not effective at using AI today, giving AI more autonomy tomorrow doesn't magically make you effective. It just gives you faster failure modes.

And once something has already been done based on them, those failures are much harder to contain.

Boundaries aren't optional: permissions, identity, traceability, containment

When an agent is connected to real systems, boundaries become the real product. Not the model. Not the UI. The boundaries. Because once agents can act, control stops being abstract and becomes the product itself.

Boundaries include:

Least-privilege permissions: the agent gets only the minimum access needed.

Credential handling: where keys and tokens live, and how they're rotated and limited.

Human approval: which actions require explicit confirmation.

Traceability: knowing what the agent did, why it did it, and what information it used.

Containment: keeping agent actions inside sandboxes where failures don't cascade.

This isn't theoretical. We have live examples of what happens when powerful agentic tools collide with real-world deployment.

In February 2026, a Reuters report said **China's industry ministry warned about security risks linked to the OpenClaw open-source AI agent**, noting that improper configurations could lead to cyberattacks and data breaches, and urging audits and stronger identity and access controls.

In January 2026, **Malwarebytes described an impersonation campaign following Clawdbot's rename to Moltbot**, including typosquat domains and a cloned GitHub repository, which is classic supply-chain risk behaviour.

Tenable published **mitigation guidance for Clawdbot/Moltbot/OpenClaw**, including **restricting token access, auditing configurations, and limiting exposure to untrusted channels** to reduce prompt injection risk.

This cluster of events matters because it illustrates what "agent risk" looks like in practice. Which is not a sentient AI plotting destruction. Instead it's a messy ecosystem of open tools, misconfiguration, supply-chain abuse, and over-permissioned systems.

That's the kind of risk organisations are actually dealing with, right now.

Autonomy works when it's built on structure

The uncomfortable truth about agents is that they're most reliable in environments that are already structured.

That means:

Workflows with clear start and end states.	Rules that define "must not happen."
Known exception paths.	Measurable outcomes.
Send the email.	Schedule the follow-up.

This idea isn't new. It's why classic automation has been valuable for decades.

What's new is adding flexible language and reasoning capabilities on top of structured automation.

The **HBR Analytic Services** report we've talked about before makes this explicit by **pointing to robotic process automation (RPA) as deterministic and rule-based**, and framing that determinism **as a natural guardrail** when layering agentic AI into real business processes.

So if you want a grounded way to think about agents, think less "digital employee" and more "automation with judgement," where judgement is tightly bounded by rules and approval points. That's where agents become useful without becoming a liability. And without that structure, agents don't simplify work.

They just scale the same problems faster.

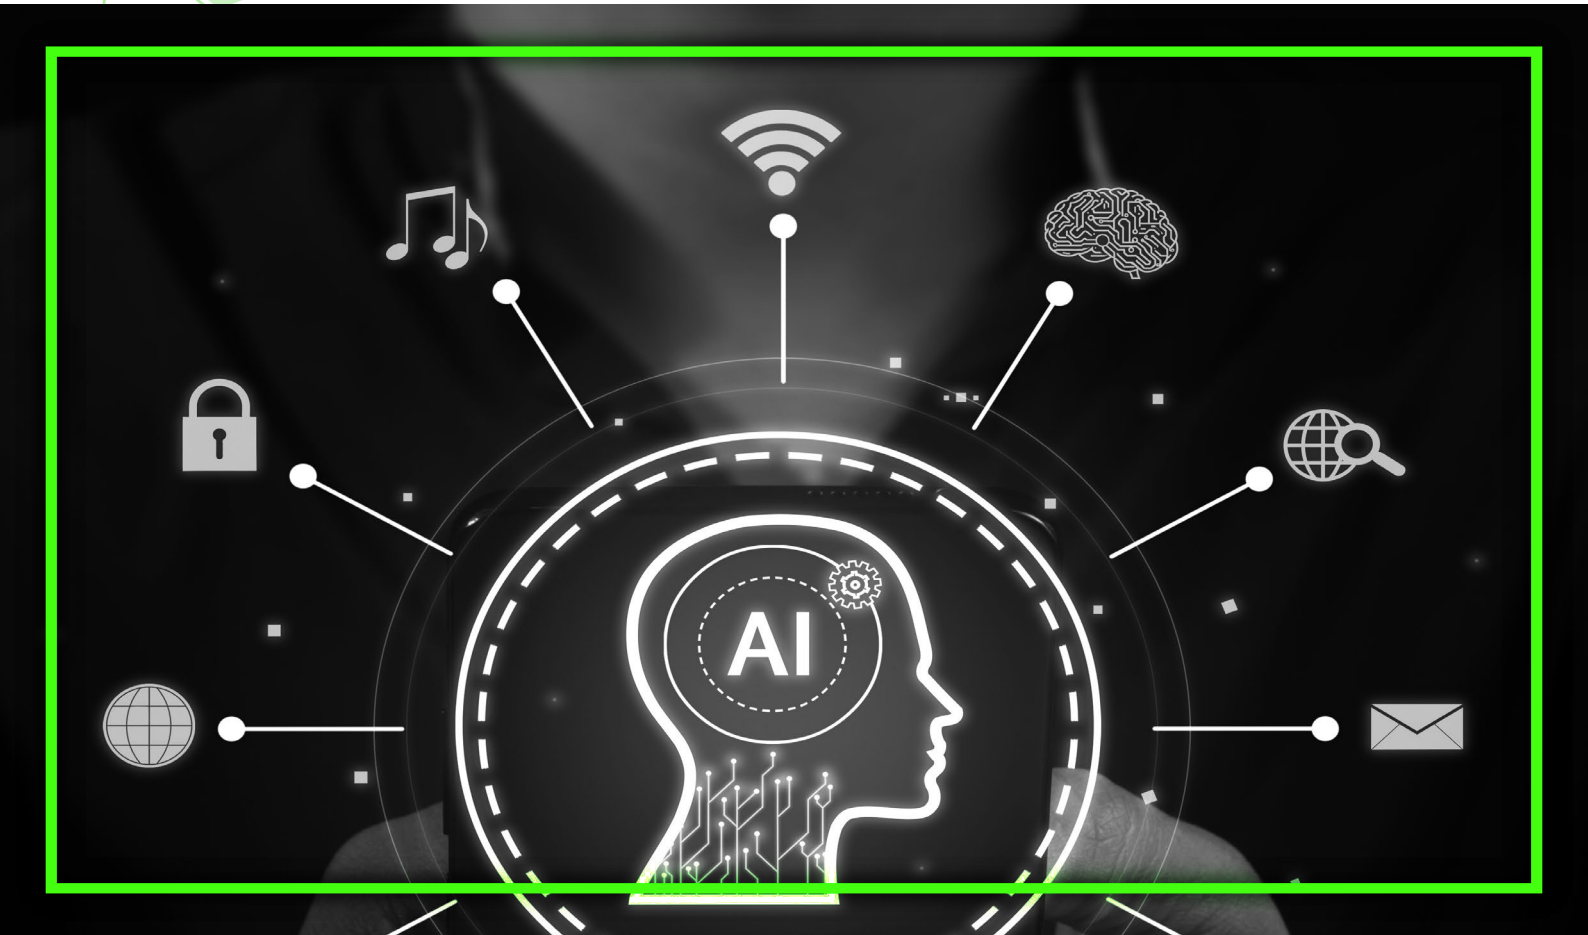

SO WHAT'S THE REAL AI STORY

It only looks contradictory up close. Step back, and it starts to make sense.

If you take a step back from all of this, the confusion starts to make a lot more sense. People aren't wrong about what they're seeing. The problem is that each signal is being treated like the whole story, instead of one part of a bigger shift that's happening underneath everything.

And when you put those pieces together, the picture that comes through is a lot less dramatic, but a lot more structural.

AI isn't collapsing. It's embedding into the systems people already use every day. That's why it can feel quieter than it did a year ago, even though adoption is still growing. It's moving out of the spotlight and into the background, where work actually happens. At the same time, it isn't replacing work in one clean sweep either. Some roles are being cut, and that part is real.

But most of the change is happening inside jobs, where specific tasks are getting faster, others are getting more complex, and expectations are shifting around both. That's why the experience feels uneven. It depends on which part of the work you're looking at.

The same pattern shows up with the idea that AI is "breaking everything." In most cases, it's not creating problems out of nowhere. It's making existing ones visible.

When outputs are wrong, when workflows slow down instead of speeding up, or when trust starts to drop, it usually points back to something that was already unclear, unowned, or stretched too thin. The technology just brings it to the surface faster, and at a scale that's harder to ignore.

And then you get to agents, which is where things start to feel like a bigger leap. But even there, it's not a separate story. It's the same pattern, just pushed further. If current systems mostly generate outputs, agents start to take action. Which means the same gaps don't just show up, they carry through into what actually gets done.

"The first rule of any technology used in a business is that automation applied to an efficient operation will magnify the efficiency. The second is that automation applied to an inefficient operation will magnify the inefficiency." ~ Bill Gates, Microsoft Co-Founder and Philanthropist

So the real story isn't hype, and it isn't collapse either. It's a shift in where AI sits and how it interacts with work. It's becoming part of the operating layer. Not something you go to, but something that sits inside the tools, the workflows, and the decisions that shape how organisations run.

And once you see it that way, the difference between what works and what doesn't becomes a lot clearer. The organisations that treat AI like infrastructure, and take the time to redesign how work actually happens around it, are the ones that will get consistent value out of it.

The ones that treat it like a tool you can add on at the edges will keep running into the same pattern. A few strong results, a lot of friction, and a growing amount of time spent checking what the system produces. That's why this moment feels messy.

Not because the direction is unclear, but because the systems around it haven't caught up yet.

WHAT THIS MEANS FOR HOW ORGANISATIONS SHOULD APPROACH AI NOW

Understanding the pattern is one thing.
Working with it properly is something else.

By this point, the pattern is pretty clear. The noise, the excitement, the frustration, even the fear, it all comes from the same place. AI is moving fast, but the way organisations are using it hasn't caught up yet.

So the question isn't really "what is AI doing?" anymore. We've already seen that. The better question is what to do about it in a way that actually works in the real world.

Because the gap isn't capability. The models are already good enough to be useful. The gap is how organisations are approaching them, where they're placing them, and what they expect to happen once they're there.

That's where most of the friction comes from. And it's also where most of the opportunity sits.

Start with the work, not the model

A lot of teams are still starting in the same place. They want to know which model is best. Which one is smarter, faster, more advanced. It sounds like a sensible question, but it usually leads nowhere useful.

What matters much more is where work is already starting to change. Not in theory, but in the day-to-day tasks people are actually doing. Writing documents, analysing data, responding to customers, preparing reports, reviewing outputs.

That's where AI is already having an effect, often quietly, inside the tools people use every day.

If you start there, the conversation becomes much clearer. You're not choosing a model in isolation. You're looking at where it fits into real workflows, where it removes friction, and where it introduces new work. That's where the value is.

Not in the model itself, but in how it's used.

Fix the system before you scale the technology

One of the biggest mistakes organisations are making is trying to scale AI before they've fixed the things around it. If your data is messy, your workflows are unclear, or no one really owns a process end to end, adding AI doesn't solve that. It just makes it more visible, and usually speeds up how quickly those issues show up.

The same applies to governance. If there isn't a clear way to check outputs, trace decisions, or step in when something goes wrong, the problem doesn't stay contained. It spreads.

That's why so many deployments feel like they stall or underdeliver. Not because the technology doesn't work, but because everything around it isn't set up to support it properly. If you want better results, you don't start by adding more AI.

You start by cleaning up the system it's going into.

Redesign jobs at the task level, not the role level

The concern around jobs is real. People are seeing roles change, and in some cases disappear, and AI is part of that story.

But most jobs aren't a single thing. They're a collection of tasks, some repetitive, some creative, some decision-based.

AI tends to affect those tasks unevenly. It might handle the first draft, but not the final judgement. It might speed up analysis, but still rely on someone to interpret the outcome.

The organisations handling this well aren't just asking which roles can be removed. They're looking more closely at what the job actually involves and how those pieces can be redesigned.

That often leads to better outcomes than simple replacement. Work gets reorganised, not just reduced.

And in many cases, the role becomes more focused on what actually requires human judgement. That doesn't remove the risk.

But it does give you more control over how change happens.

Treat AI output as a starting point, not a finished product

One of the easiest traps to fall into is trusting the output because it looks right.

AI is very good at producing something that feels complete. The language is polished, the structure makes sense, and the answer arrives quickly. That creates a natural tendency to move on without checking too closely.

But that's where the hidden work starts to build up.

There's a clear gap between how ready organisations think they are and how ready they actually are. Deloitte's 2026 State of AI in the Enterprise report found that 42% of companies say their strategy is highly prepared for AI adoption, but far fewer feel ready when it comes to infrastructure, data, risk, and talent.

Someone still has to review it, validate it, and decide whether it can be used. And as usage grows, that review burden grows with it.

So if that part isn't made explicit, it turns into a silent cost.

People spend more time checking, correcting, and second-guessing, often without it being recognised as part of the job.

Treating AI output as a starting point keeps that visible. It sets the expectation that the work isn't finished when the answer appears. It's finished when it's been properly understood and checked.

Define boundaries before you increase autonomy

As soon as AI starts taking action instead of just producing output, the stakes change. Now it's not just about whether something is right or wrong. It's about what happens if it acts on something that isn't.

That's why boundaries matter. What the system is allowed to access, what it's allowed to do, where it needs approval, and where a human needs to step in. Without that, you still get decisions and actions, but with far less clarity and control over how they happened.

And once something has already been done, it's much harder to contain the impact. If you want to use more autonomous systems, those boundaries need to be defined first. **Not added afterwards.**

Assume work will feel more complex before it feels easier

There's an expectation that faster tools should make work feel simpler. In practice, that's not what tends to happen, at least not at first.

When things move faster, expectations increase. More gets produced, more needs to be reviewed, and more decisions happen in less time. Work doesn't disappear. It expands, and it often becomes more layered. That can feel like something is going wrong. Like the tools aren't delivering on the promise.

But in many cases, it's just the reality of introducing something powerful into an environment that hasn't fully adjusted to it yet.

Planning for that makes a difference. If you expect work to feel more complex for a while, you can design around it. If you expect immediate simplicity, it will always feel like a disappointment. Over time, things do settle. **But only if the way work is structured catches up with the tools being used.**

FINAL THOUGHTS: THE REAL AI STORY IS ABOUT HOW WE LEARN TO LIVE WITH IT

It doesn't become manageable by default. It becomes manageable by design.

If you go back to that feeling we started with, the confusion, the back and forth, the sense that nothing quite lines up, it makes more sense now.

You're not looking at one clean story unfolding. You're looking at multiple layers of change happening at once, and not all of them moving at the same pace.

The loud phase trained everyone to focus on moments. New models, big claims, strong reactions. But most of what actually matters now isn't happening in those moments. It's happening in the background, inside the systems people already use, where work gets done and decisions take shape.

That's why it feels harder to pin down. Not because progress has stopped, but because it's become part of how things already work.

And once that happens, the conversation stops being about what AI can do in isolation. It becomes about what happens when that capability sits inside real organisations, with all the complexity that comes with them.

That's where the pressure shows up.

Not as one obvious failure, but as a pattern that repeats in different forms. Things move faster, but expectations rise with them. Outputs improv

e, but trust becomes harder to manage. Some parts of work get easier, while others quietly become heavier.

When you look across all of it, the same issue keeps coming up. The gap isn't capability. The systems are already good enough to be useful in a lot of places. The gap is in how organisations are set up to work with them.

That's what readiness really comes down to.

It's about whether workflows are clear enough to support something that moves this quickly, whether ownership is defined when outputs are generated across systems, and whether people know when to rely on the system and when to step in.

It's about setting boundaries before more autonomy is introduced, not trying to recover control after the fact. And it's about recognising that these systems don't run in isolation — they depend on infrastructure, cost, and constraints that shape what's actually sustainable.

That's also why the experience feels so uneven right now. Some teams are getting real value because they've adjusted how work happens around the technology. Others are running into friction because they've added it on top of processes that weren't designed for it.

So if you want one steady takeaway, it's this:

The difference between what works and what doesn't isn't coming from the technology itself. It's coming from how deliberately organisations adapt the way they work around it.

That's where the real story sits now.

Not in what AI might be able to do next, but in how people, teams, and organisations learn to work with something that's already here.

EM360Tech

connects enterprise technology leaders with the insights, analysis, and conversations shaping the future of AI, data, security, and infrastructure.

For more insights,
visit www.em360tech.com
or contact us to collaborate on a thought leadership campaign:
info@em360tech.com