


The Only Human in the Room

What AI agents mean for any organization
whose product is information



July 2026

The Only Human in the Room

By [Mohammed Juzer Batliwala](#) and [Andrej Verity](#) | [DigitalHumanitarians.com](#)

Key messages

- **People started without permission.** Most employees already use AI at work, often without telling anyone, and only about four in ten workplaces have any rules about it. The bottleneck is not the technology; it is how organizations are set up.
- **Your readers are becoming machines.** During 2025, the share of website visits made by AI tools rather than people rose from roughly *one-in-two-hundred* to about *one-in-thirty*. If your product is information, a growing part of your audience is software.
- **AI tools do not always read what they cite.** In our test, one AI browser read the live page in ten of twelve questions. Another read it in none, while still showing a citation to the page it never opened. Your name can end up on an answer you never gave.
- **Get ready fast. Hand over control slowly.** Let AI in where mistakes are easy to catch and fix. Keep people in charge of the decisions that cannot be undone.

The gap runs in two directions

Most organizations are having the wrong argument about AI. The debate gets framed as whether to adopt it or not. The evidence says adoption already happened, without anyone asking, and the real question is what to do about it.

Two things are true at the same time. Staff are ahead of their organizations: in the largest global study of the question, spanning 47 countries and more than 48,000 people, a majority of employees said they use AI at work in ways their organization has not approved, more than half admitted to presenting AI-generated work as their own, and only about four in ten workplaces had any policy on it at all.¹ A separate study found employees at more than 90

¹ Gillespie, N., Lockey, S., Ward, T., Macdade, A., and Hassed, G. (2025). *Trust, Attitudes and Use of Artificial Intelligence: A Global Study 2025*. University of Melbourne and KPMG. Survey of 48,340 people in 47 countries.

percent of companies using personal AI tools while only about 40 percent of those companies had bought an official subscription. The researchers called it the shadow AI economy.²

And the technology is behind its own marketing. When AI agents are tested on realistic office tasks, even the best complete only about a quarter to a third of them, and they sometimes present a failed task as a success.³ Gartner, whose forecasts drive much of the enthusiasm, at the same time predicts that more than 40 percent of agentic AI projects will be cancelled by the end of 2027, and estimates that of the thousands of vendors advertising AI-agent products, only around 130 are the real thing.⁴

The honest starting point is that people are already using AI everywhere, and the tools are not yet as good as advertised. Neither racing to catch up with the hype nor waving it away is a sensible response.

Your readers are becoming machines

Whatever you make of the forecasts, the plumbing is real. The Model Context Protocol, an open standard that lets AI tools connect to data and services, went from launch to more than ninety-seven million downloads a month by developers within a year, and was handed to the Linux Foundation in December 2025 so no single company controls it.⁵ Google's equivalent standard for agent-to-agent communication followed the same path. Meanwhile, measurements of open-web traffic found the share of visits made by AI tools rather than people collapsing from roughly one-in-two-hundred at the start of 2025 to about one-in-thirty-one by the end of the year, with the fastest growth in tools that fetch a page to answer somebody's live question.⁶

If your product is information, this changes who your audience is. The visitor you are designing for is, increasingly, a machine acting for a person.

² Challapally, A., Pease, C., Raskar, R., and Chari, P. (2025). *The GenAI Divide: State of AI in Business 2025*. MIT Project NADA.

³ Xu, F. F., and colleagues (2024). *TheAgentCompany: Benchmarking LLM Agents on Consequential Real World Tasks*. arXiv:2412.14161.

⁴ Gartner press release, 25 June 2025, and Gartner planning assumptions published 2025 to 2026.

⁵ Linux Foundation announcement, December 2025: launch of the Agentic AI Foundation and contribution of the Model Context Protocol.

⁶ Tollbit (2025). *State of the Bots, Q4 2025*.

What we found

We ran a structured comparison of three consumer AI browsers, pointing each at the same authoritative source with the same twelve questions, and scoring how they answered.⁷

They behaved very differently. One read the live page in ten of twelve questions. One managed six. One read it in none, answering instead from its own search index or training memory in most cases, while still attaching a citation that pointed at the source it had not opened.

Two lessons follow. First, accuracy depended on whether the tool actually read the source, not on how hard the question was. That is the one thing an information provider can influence. Second, and more troubling: your organization's name can end up attached to an answer you never gave. That is a problem of trust in what you are quoted as saying, not a problem of server load, and blocking bots does not fix it.

Where AI belongs

A simple rule: ask what happens if the AI gets something wrong, and whether the mistake can be undone.

- **AI earns its keep in the middle ground:** drafting documents, pulling information together, monitoring, translating. A mistake there gets caught and fixed by whoever reviews the work.
- **At the trivial end it just adds friction:** an extra step for something a person does faster alone.
- **The high-stakes calls stay with people:** anything that cannot be taken back, or where someone outside the organization will hold you to account. The AI can prepare the work. A person makes the call.

What to do

Move fast on getting ready. Move slowly on handing over control.

⁷ Author's structured comparison of three consumer AI browsers: twelve identical questions per browser, scored on accuracy, sourcing, recency, completeness, and fabrication, May 2026. Preliminary; single session.

Do these five things now:

1. **Publish official interfaces** that let AI tools read your data properly.
2. **Describe and label your data in ways machines can read.** Fund it as ordinary data-quality work, not as an AI project.
3. **Measure AI traffic** the way you measure human visitors, starting by telling apart bots that copy everything from tools answering someone's live question. Almost nobody can make that distinction today.
4. **Write a clear access policy for AI crawlers** to replace rules written before ChatGPT existed.
5. **Take a seat in the groups now writing the technical standards**, so your requirements are built in from the start.

Why now: the industry expects these standards to be settled around 2027 and 2028.⁸ A standard is easy to influence while it is being written and nearly impossible to change afterwards. Wait, and you inherit rules designed by and for someone else.

And hold the line here: do not let AI make the calls that cannot be undone. Every technical layer of the agent economy is being built: identity, payments, communication between agents. The one layer nobody has built is the one that matters most: who is responsible when an autonomous system makes a mistake.

The risk nobody is pricing

The price of running AI today is artificially low. It is held down by free credits, below-cost pricing, and compute that loses money for the companies selling it. If you build on a vendor's discount, you have a hidden dependency that only shows up when the discount ends. Favor open, portable standards for this reason as much as any other, set cost alerts on your AI usage, and plan for vendors to disappear rather than assuming they will still be there.

One more thing, on the human side

The productivity research is encouraging. A field experiment with 776 professionals found that individuals working with AI matched the output of two-person teams working without it,

⁸ Gartner press release, 25 June 2025, and Gartner planning assumptions published 2025 to 2026.

and that AI dissolved the walls between specialisms.⁹ But a study across four countries found that the more employees worked with AI, the lonelier they reported feeling, with effects that followed them home into sleeplessness and after-work drinking.¹⁰

These findings do not contradict each other. Enjoying a task and feeling connected to people are different things. An AI can make the hour more engaging while removing the colleague you would have shared it with. The work a teammate does can be replaced. The company a teammate provides cannot. Organizations that measure only the first will hollow out the second without noticing.

To be the only human in a room full of AI is not to be obsolete. It is to hold the position the machines cannot: the one accountable, in the end, for what the work is for.

⁹ Dell'Acqua, F., and colleagues (2025). *The Cybernetic Teammate*. NBER Working Paper 33641.

¹⁰ Tang, P. M., and colleagues (2023). No Person Is an Island. *Journal of Applied Psychology*, 108(11), 1766–1789.