




The
AGENTIC
STATE of **AI**

FIELD NOTES FROM Q3 2026



A note before you read

Over the past few months, I've gotten to speak with more than 50 media outlets and many more people at companies spending heavily on AI and trying to work out what they're actually getting for it. After enough of those conversations, I kept hearing the same thing under the questions. While there's plenty of fear and plenty of hype, the consistent theme was the space between what organizations hope AI can do and what's actually happening inside their (digital) walls.

I've started calling it the gap between ambition and execution. The strange thing about this moment is that the gap is showing up everywhere at the same time. Earlier waves of tech, like the PC, tended to hit one industry first and spread from there. But this one arrived in healthcare, government, supply chains, and marketing all at once, so almost everyone I talk to is working through some version of the same problem. It's a zeitgeist moment, which isn't something you get to say very often.

I didn't write this as a how-to. I wrote it to capture what I'm seeing in Q3 2026: **five tensions playing out between what the market believes and what the data shows, along with what organizations making the most progress are doing about each one.** I'm approaching this from the data layer, because that's the one place where the shadow usage, the governance gaps, and the financial impact can come together in the same view. It gives use a useful way to see how these pieces connect.

If you'd rather hear this than read it, the full conversation that this guide came out of is up as a [video podcast with Mark Boothe, Domo's CMO](#). Otherwise, read on.



*Chris Willis
Chief Design Officer and
AI Futurist, Domo*

CONTENTS

Tension 1 4

Productivity gets worse before it gets better

Tension 2 5

Most AI usage is invisible to leadership

Tension 3 6

The order that separates the 5% from the 95%

Tension 4 7

Confidence is outrunning control

Tension 5 8

The trust problem and the VPC model

Closing note 9



TENSION 1

Productivity gets worse before it gets better

By now, AI is supposed to be producing a measurable boom. The money went in, so output should be climbing.

The data shows a dip instead, and economists have seen this shape before. When a new general-purpose technology arrives, measured productivity usually falls before it rises. The curve traces a J: a baseline, a drop, then a long climb.

The personal computer (PC) is a perfect example. In the late 1980s and early 1990s, PCs landed on desks, and most people had no idea what to do with them. Which way the floppy disk went in wasn't obvious. And at first no one knew what to do with the mouse at first.

~50%

THE DROP IN MEASURED OFFICE PRODUCTIVITY ACROSS THE FIRST DECADE OF THE PC, BEFORE THE PAYOFF ARRIVED ABOUT TEN YEARS IN*

The dip is mostly invisible investment that doesn't show up as a budget line item:

- Learning the tools well enough to trust them.
- Rebuilding the workflow around them, usually getting it wrong a few times first.
- Working out which parts of a job the technology is actually good at, and which parts it only looks good at.

It's slow, and it costs time. Agents are creating this same uncounted work right now, which is why so many teams feel busier and slower in the same week.

The pattern tends to move in three phases:

- The technology arrives and the excitement peaks.
- Productivity dips while people build the skills, workflows, and invisible infrastructure around it.
- The returns start to compound.

That old line about the future already being here but not evenly distributed fits the moment. A few organizations are already seeing returns; most are still working through the dip. The organizations we see getting ahead treat that period as something to manage deliberately rather than as a warning to turn back. Much of the rest of this guide is about what helps them move through it faster.



TENSION 2

Most AI usage is invisible to leadership

Most leaders believe they can see how AI is being used across the company. They point to the sanctioned pilots and the official numbers and assume that's the picture.

But it's the tip of an iceberg. The pilots are real, but they float on a much larger layer of everyday usage that never gets reported. Most of what people actually do with AI happens below the waterline:

- Above the waterline: Sanctioned pilots and dashboards leadership actually reviews.
- Below the waterline: Personal tools and side workflows that aren't tracked or logged with IT.

The numbers show how wide that gap can be.

~5%

OF SANCTIONED, TOP-DOWN AI PILOTS SHOW A MEASURABLE P&L RETURN

90%+

OF FIRMS ALREADY HAVE EMPLOYEES USING PERSONAL AI TOOLS

3x

HOW MUCH LEADERS UNDERESTIMATE AI USE AMONG THEIR PEOPLE*

Underneath this lies a coordination problem that predates AI. People find small efficiencies on their own, but the structure that would spread those wins across the organization usually isn't there.

Call it the dysfunction of hierarchy: a standing gap between what leaders believe is happening and what the front line lives with. Agents threaten to blow it open.

Agents make that problem much larger. A team of 10 people is still 10 people, and you can more or less keep track of them. 10 people each running a thousand agents is ten thousand loose threads moving at once, most of them out of view.

That same blind spot can lead organizations to automate too aggressively. One widely covered company bet big on automating its sales-development reps out of existence, and not long after, its CEO announced they were hiring human SDRs again.

The technology could handle pieces of the job, but it had never seen their business or the judgment calls that came with the accounts. The teams finding their footing tend to use AI for repetitive work while keeping people involved where context and judgment is needed.

The leaders we see getting value first make the hidden usage visible before they scale. They read the tools their people already reach for as a signal about where the real demand is. Trying to stamp that usage out mostly pushes it further underground, and none of the coordination works while it stays hidden.



TENSION 3

The order that separates the 5% from the 95%

The instinct is to buy your way in. More tools should mean more innovation, so the plan becomes purchase the capability, mandate adoption, and wait for the results to land.

But that tools-first approach causes a lot of teams to stall. If you hand people a stack of new tools and ask them to become innovators, your spend is going to go up, but not your productivity.

92%

OF COMPANIES ARE RAISING THEIR AI SPEND

95%

OF PILOTS STILL SHOW NO MEASURABLE P&L IMPACT

~5%

OF ORGANIZATIONS ARE ACTUALLY READY TO USE AI, BY MCKINSEY'S MEASURE*

Part of the problem is what gets measured. When the metric is tool usage, you get what people started calling tokenmaxxing: dashboards full of adoption numbers with no outcome attached. Employees end up performing innovation instead of doing their actual work, most of them pushed well outside what they're comfortable doing.

The order that tends to work runs the other way around:

- **Mindset first.** Start from the decision you're trying to improve, get curious about it, and find the people already working on it. Culture is the operating system of the organization, and that's the layer that has to shift.
- **Enablement next.** Give those ideas the data, workflows, governance, and training they need to spread without creating a mess. This is much of the invisible infrastructure behind the productivity dip.
- **Tools last.** Choose ones that fit the workflow you now understand instead of asking the workflow to bend around whatever was purchased.

This approach also shifts who you hire. The work rewards generalists with range, people who move across disciplines and stay comfortable with tools they don't fully understand yet. I asked the head of a Fortune 100 company recently how AI was changing hiring, and they said they were looking for people with better writing skills. The theory was that good writers tend to be good thinkers, and since the AI models work well with language, this skill has become increasingly valuable.

The other requirement is humility. When an expensive bet flops, the teams that recover ask which specific part flopped and why, rather than burying it. That habit is how an organization learns fast enough to cross the dip.



TENSION 4

Confidence is outrunning control

Ask an executive whether their AI agents are under control and most will say yes. The policies are written, the guardrails are in place, and the assumption is that agents stay inside them.

That confidence is running well ahead of what many companies can see and govern. At least 82 percent of executives are confident their policies prevent agents from taking unauthorized actions.

However:

ONLY

21%

PERCENT HAVE A MATURE MODEL FOR GOVERNING AGENTS OVER

50%

OF AGENTS RUN WITH NO SECURITY OVERSIGHT OR LOGGING AT ALL

88%

OF ORGANIZATIONS REPORTED A CONFIRMED OR SUSPECTED AGENT SECURITY INCIDENT IN THE PAST YEAR*

The gap is hard to miss. Most leaders trust the policy, but far fewer can actually see what their agents are doing, and a large majority have already had something go wrong.

This is what blocks agents from reaching production next. The invisible infrastructure that grew out of the productivity dip. It also includes the parts that make a fleet of agents safe to run and easier to govern:

- Monitoring, so someone can see what an agent actually did.
- Logging, so there's a record when something goes sideways.
- A governed model of data, so agents reach the information they should and not the information they shouldn't.
- A clear owner, so accountability doesn't evaporate.

Accountability is easy to overlook. But when an agent spins up a new ad campaign or a landing page and it goes wrong, someone has to answer for it. 'The model did it' does not hold up in front of a board.

The organizations we see moving fastest are the ones closing the gap between confidence and control before they scale. They build in governance and visibility early enough to understand how agents are behaving once they reach production.



TENSION 5

The trust problem and the VPC model

The stated goal, when people talk about deploying agents, is usually to make the AI trustworthy by getting the models reliable enough. A reliable model equals a trustworthy model, so the thinking goes

But trust doesn't work that way in real workplaces. Trust is a human judgment, something people give or withhold, and you can't order it into place or wire it straight into the model. The system around the model, though, is something you can build.

66%

USE AI REGULARLY

46%

ACTUALLY TRUST IT

56%

REPORT MISTAKES AT WORK CAUSED BY AI

1 IN 3

ORGANIZATIONS HAVE REACHED A MATURE LEVEL OF AI TRUST*

Building trust comes down to three things: visibility, predictability, and control.

Visibility. You can see what the agent actually did, step by step: the real actions it took and the data it touched, rather than relying on the tidy summary it writes about itself afterward. In the projects we're building at [Domo](#), this goes straight into the interface. When an agent recommends something, it shows why, how, and which sources it used.

Predictability. Agents aren't predictable, and they don't always approach the same situation with the same behavior; there's randomness to it. But you can narrow the room they have to improvise. One way is certified data products: Instead of letting an agent write a fresh SQL query every time it needs revenue data, you can point it to a governed revenue metric with fixed parameters, so the number comes back safe and repeatable. Any behavior you haven't defined is one the agent will invent, so the goal is to keep that space small.

Control. You need have clear limits on what an agent can do and the ability to stop it outright, with a human in the loop and guardrails that hold even when the model pushes against them. Before an agent edits a JIRA ticket on your behalf, for instance, it should have to ask you to approve the change first.

These characteristics are the difference between an agent that stays a pilot and one you can put into production. Until you can see what it does and rein it in, it stays an interesting experiment. Once you can, it becomes something you can give responsibility.



How the 5 tensions resolve

Every tension in this guide comes back to the same requirement: **an integrated, governed data environment**. When the data is spread across places like a warehouse, an ETL system, a visualization tool, and other disconnected parts of a kind of Frankenstein stack, you can't see the whole picture. This splintering is why most of these gaps persist. A lot of organizations are running a kind of Frankenstein stack.

A system built around the data layer makes that activity visible and gives organizations the foundation for visibility, predictability, and control. When we've said Domo was made for this moment, this is what we meant. We built around the data layer without knowing this moment was coming, and it to be exactly where so many of these challenges converge. Sometimes it's just better to be lucky.

Which leaves the question I keep coming back to. In cosmology, there's the idea of a great filter, something that helps explain why we look out at the universe and don't find the intelligent life we might expect.

There's a version of it here. If everyone has access to the same AI, why are some organization getting much further with it than others? My guess is that the ones through the filter first are the ones that can scan across their whole organization and learn from it faster than anyone else.

Chris Willis
Chief Design Officer and
AI Futurist, Domo

