



Shadow AI Isn't Your Problem. It's Your Operating Model Telling You the Truth.

Cheney Hamilton ©Bloor Research October 2026

Why the productivity gains attributed to Artificial Intelligence are, in many organisations, being subsidised by work nobody has written down

Officially, the transformation programme is on track.

The dashboard is green. The workstreams are reporting. The steering group met last week and nobody raised anything that could not be managed.

Underneath, something else is happening.

Work is being coordinated in messaging groups outside approved platforms. Managers are exporting data into private spreadsheets because the core systems cannot produce what they need. And employees are using personal Artificial Intelligence (AI) tools to finish tasks faster than the sanctioned systems allow.

Most organisations treat this as a compliance problem.

I think it is better read as a diagnostic.

Because the gap between those two realities — the reported one and the actual one — is not a governance failure. It is information. And it is some of the most accurate information your organisation currently produces about how work really gets done.

What Shadow AI actually is

The term has arrived quickly and it is already being used loosely, so it is worth being precise.

Shadow AI is the use of unsanctioned AI tools to perform work that the organisation's approved systems either cannot do, or cannot do fast enough.

It is not usually malicious. It is rarely even covert in any meaningful sense — ask most teams directly and they will tell you exactly what they are using and why.

It is people doing their jobs with the best tools available to them, which happen not to be the tools they were given.

And it is spreading considerably faster than most governance functions have registered, because the barrier to entry is a browser tab and the benefit is immediate.

There is a detail here that organisations consistently miss. The people most likely to be doing it are not the disengaged ones. They are the

capable, motivated, time-pressured people who care about getting the work out — which is to say, the people you can least afford to alienate with a blanket ban.

The instinctive response is to police it. Block the domains, issue the policy, run the training, add it to the risk register.

I would suggest that is the wrong first move, because it removes the evidence without addressing what the evidence is telling you.

The workaround network is your real operating model

Here is the uncomfortable observation underneath all of this.

Most organisations do not run on their documented operating model.

They run on human workaround networks that emerged because the official systems stopped reflecting how work actually happens.

This is not new. The spreadsheet that sits beside the enterprise system, the informal escalation route that bypasses three approval stages, the person everyone rings because they know how to make the process work — these predate generative AI by decades.

What is new is the acceleration.

Generative AI did not create those workarounds. But it has dramatically increased both their scale and their invisibility. A workaround that used to be constrained by how fast a human could type is no longer constrained by that at all.

And here is why none of it surfaces: the workarounds are invisible precisely because they work. Nobody escalates a problem they have already solved. The spreadsheet gets built, the data gets reconciled, the deadline gets met — and the system that failed to support any of it continues to report green, because from where the dashboard is standing, nothing went wrong.

Gartner's work on operating-model redesign frames the same shift from the enterprise-architecture side: the unit of design is moving from the job to the outcome.

Which is another way of saying that the documented model — built around jobs, roles and defined processes — has stopped describing the thing it was drawn to describe.




Figure 1 – The operating model gap. What the dashboards see, and what actually runs the organisation

The subsidy nobody has costed

Now for the part that should concern any Chief Operating Officer (COO) signing off a business case.

In many organisations, the productivity gains being attributed to AI are being subsidised by undocumented human intervention.

Someone is checking the output before it reaches a client. Someone is rewriting the prompt because the first three attempts were not usable. Someone is reconciling the two systems that still do not talk to each other. Someone is quietly correcting the thing the model got confidently wrong.

None of that work appears in the business case.

Take a process reported as forty per cent faster. Plausible, and probably true at the point of execution. But if the time saved reappears as a senior person spending two hours a week quality-checking output that was not reliable enough to send, the net position is considerably less impressive — and the cost has moved from a junior grade to an expensive one.

Which means the business case is wrong — not catastrophically, not deliberately, but systematically and in one direction.

The gains are reported in full. The costs of sustaining them are absorbed invisibly, by people whose job descriptions do not mention any of it and whose managers may not know it is happening.

If you want one question to take into your next AI review, I would make it this one:

what human work is making this number possible, and who is doing it?

In my experience, nobody in the room can answer it immediately. That is itself the finding.

Concealment, not elimination

There is a broader pattern here that deserves naming.

AI is frequently not replacing operational complexity. It is concealing it.

Complexity that used to be visible as delay, as backlog, as an obviously overloaded team, now gets absorbed — partly by the technology, partly by people working around the technology. The symptom disappears. The condition does not.

And a concealed problem is considerably more dangerous than a visible one, because it stops generating the signals that would have prompted somebody to fix it.

This is the point at which the stakes rise. As deployment moves from task-level assistance towards agentic systems that plan and execute whole workflows, the undocumented human layer holding things together becomes both more necessary and harder to see.

Read the signal, don't police it

None of this is an argument for abandoning governance. Unsanctioned tools carry real risks around data, confidentiality and accuracy, and those risks need managing.

But governance alone treats the symptom.

If people are routing around your systems at scale, the useful question is not how to stop them. It is what the routing is telling you about the gap between your documented operating model and the one that actually runs the place.

In practice that means three things. Find out what people are actually using, without penalty attached to the answer — you will not get an honest map any other way. Treat each workaround as a specification: it describes a capability your sanctioned systems do not provide. And cost the human layer propping up your AI gains before you report those gains again.



Shadow AI is not the problem. It is evidence of a hidden operating model.

Map the work at outcome level and that hidden model becomes visible — including the parts of it your organisation has been quietly depending on without ever accounting for them.

If you would like to see what your own operating model actually looks like, we are running Workforce X-Ray pilots in single functions — designed to surface outcome-level work, including the undocumented layer holding it together. If that is a conversation worth having, ***I'd like to talk to you.***

This is the fourth article in a series drawn from the Bloor Research white paper ***FusionWork™: Defining a New Category of Work.***

ARTICLE REPRINT



**BLOOR
RESEARCH**

+44 (0)1494 311 460
info@Bloorresearch.com
www.Bloorresearch.com