

Human-in-the-Loop or AI-in-the-Loop? You Can't Choose Until You Know What the Work Is.

Cheney Hamilton ©Bloor Research November 2026

ARTICLE REPRINT

The biggest design decision in your AI programme is the one most organisations never make. They inherit it from their vendor.

A choice most organisations inherit

There is a decision sitting inside every Artificial Intelligence (AI) programme that almost nobody makes on purpose.

It isn't which tool to buy, which model to use or which vendor to trust. It is simpler and much bigger than that. When a human and a digital system work on the same outcome, who leads?

Most organisations never answer it. They inherit the answer. The vendor's product was built around one way of working, the implementation partner configured it that way, and by the time anyone asks the question, the operating model has already been decided by a procurement process.

That is a strange way to make the biggest design decision you will make this decade.

Earlier in this series I argued that approximately 90% of job descriptions no longer describe the job. This is where that number stops being a Human Resources (HR) problem and becomes an operating problem. You cannot decide how humans and AI should share work you have not mapped.

Human-in-the-Loop, defined

At one end of the spectrum sits Human-in-the-Loop.

Here, digital systems do most of the execution. The human role concentrates on oversight, validation, exception handling, judgement and accountability. The machine runs the process. The person makes sure it runs properly, and owns the outcome when it doesn't.

Take accounts payable. Invoices arrive and a digital system reads them, matches them to purchase orders, codes them and queues them for payment. A finance professional reviews the exceptions: the mismatched amount, the new supplier, the invoice that looks almost right. They approve the payment run, and they are accountable for it.

Done well, this is powerful. High-volume, rules-heavy work moves faster and more consistently than people alone can manage, and human attention goes where it is actually needed.

Done badly, it is a rubber stamp. The human signs off what the system decided, without the time, context or authority to challenge it. That is not oversight. It is liability with a name attached.

The difference between those two versions is not the technology. It is whether anyone designed the human part of the loop.

AI-in-the-Loop, defined

At the other end sits AI-in-the-Loop.

Here, humans remain the primary actors. Digital capability augments, accelerates and extends what they do. The person leads. The AI amplifies.

Take an HR business partner working through a complex restructure. AI drafts the consultation documents, models the options, pulls precedent from previous cases and flags where the proposal conflicts with policy. The judgement, the conversations and the decision stay with the person, who now has more time for exactly those things.

This is the model for work where judgement, relationships and creativity carry the value. It is also where the invisible third of the work lives: the coaching, the stakeholder handling, the knowing who to call.

The evidence backs the distinction. A review of 106 experiments by researchers at the Massachusetts Institute of Technology (MIT), published in Nature Human Behaviour, found that human and AI combinations did not, on average, beat the best of humans or AI working alone. They did worse on decision tasks and better on creative ones.

In other words, the configuration matters as much as the technology. Put the human and the AI together the wrong way round and you can end up with something worse than either.



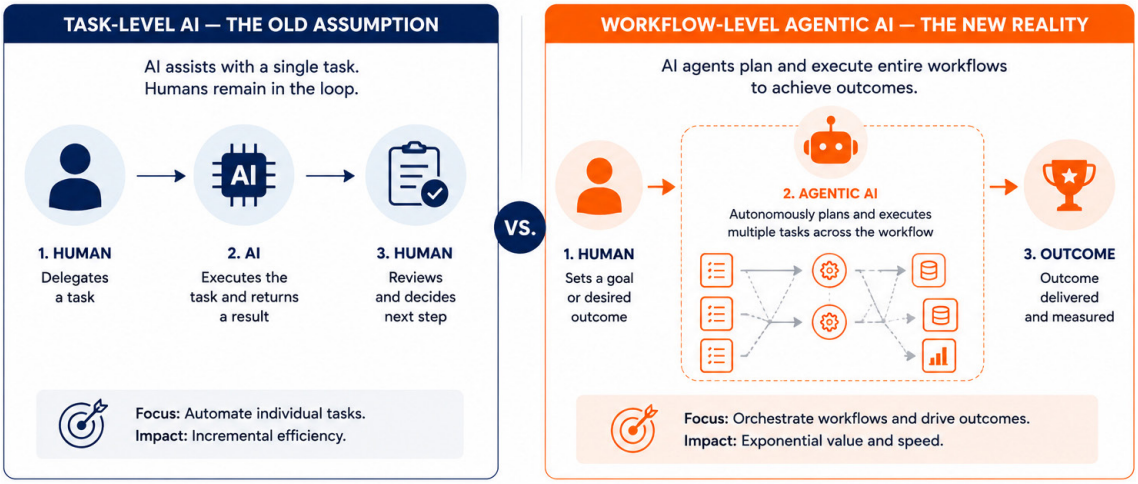


Figure 1 – The agentic shift. Task-level AI executes what a human delegates. Workflow-level agentic AI plans and delivers an outcome across a whole workflow.

The agentic shift raises the stakes

If this were still about chatbots and copilots, you could probably get away with inheriting the answer. It isn't.

The old assumption was task-level AI. A human delegates a task, the AI executes it, and the human reviews the result. The person is still the one stitching tasks together into work.

The new reality is workflow-level agentic AI. A human sets a goal or an outcome. AI agents plan and carry out multiple tasks across a workflow, hand work between themselves, and deliver an outcome that gets measured. The stitching moves to the machine.

That changes the economics. Automating tasks gives you incremental efficiency. Orchestrating workflows gives you a step change in value and speed.

It also changes the risk, because the decision about who leads is no longer made task by task. It is made once, for the whole workflow, and then it scales.

Here is the problem for a role-based organisation. Agents work across workflows. Job descriptions describe roles. If your map of work is organised by role, and most of those roles are described wrongly, you are about to hand agents a workflow nobody has actually drawn.

That is how organisations end up with Human-in-the-Loop by accident: agents running work end to end, with a person nominally accountable who was never designed into the process.

Neither can be chosen without the map

Neither end of the spectrum is wrong. Most organisations will run both, plus a range of hybrids in between, often inside the same function.

What matters is that the choice is made deliberately, from an accurate picture of the work, rather than inherited from assumptions about roles.

That means knowing, workflow by workflow, which work is judgement and which is volume. Which decisions carry legal or reputational weight. Where the invisible work sits, so you don't automate it away by mistake. And where accountability needs to land when something goes wrong.

It matters beyond your own walls, too. Earlier in this series we looked at what digital labour does to the tax base and to employment law. The choice between these two models decides how much of that exposure you create. You cannot govern a choice for tax, liability or accountability if you never actually made it.

RADBOT™: a capability cycle, not a hiring transaction

So what replaces the inherited default?

For most of the last century, organisations got capability one way: they hired it. Write a role, recruit a person, manage them, replace them when they leave. A transaction, repeated.

ARTICLE REPRINT



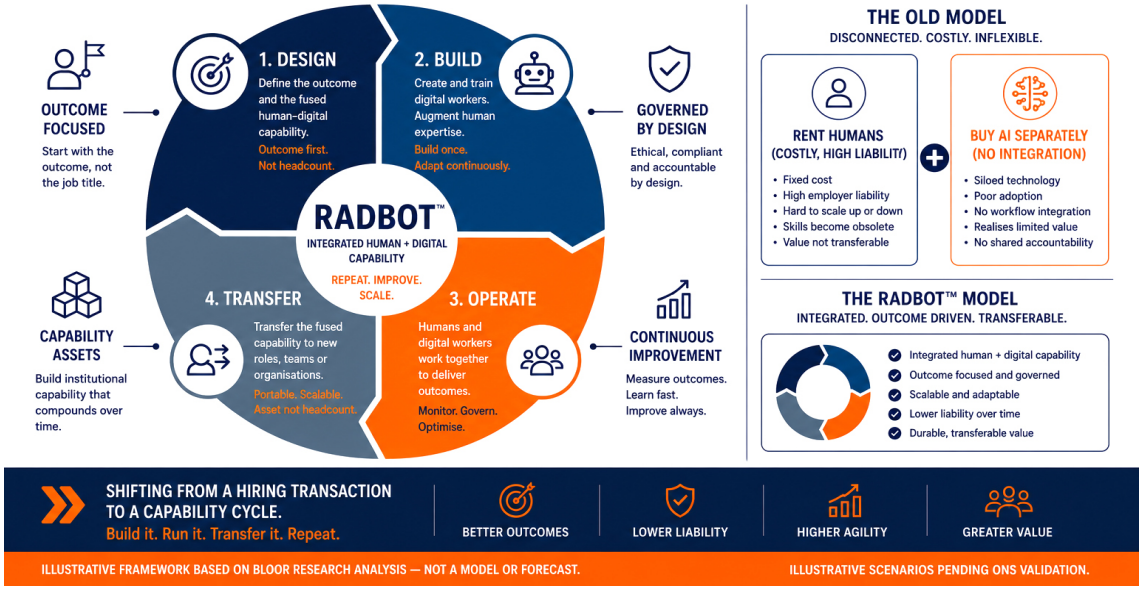


Figure 2 – The RADBOT™ cycle, old model vs RADBOT™ model. Illustrative framework based on Bloor Research analysis, not a model or forecast. Illustrative scenarios pending validation.

When capability can be human, digital or both, that transaction stops working. The FusionWork™ white paper sets out RADBOT™, a named operating cycle that replaces it. It runs in four stages.

Design the work around the outcome, not the role.

Build the fused human and digital configuration that delivers it, deciding deliberately where the human leads and where the AI does.

Operate it with governance, measurement and clear accountability, so that every automated decision has a named human behind it.

Transfer capability and value on purpose, including how a Digital Me is owned, licensed and wound down when someone moves on.

Design, build, operate, transfer. A capability cycle, not a hiring transaction. And the Human-in-the-Loop or AI-in-the-Loop decision sits right at its centre, in the Build stage, where it belongs.

Start with the map

The binary choice is not a technology decision. It is a design decision. And it cannot be made responsibly, or governed for tax, liability and accountability, without first mapping what work is actually being done.

So pick one workflow that is about to get an agent. Before anyone configures anything, map it: what the work really is, who does it, and which parts carry judgement. Then make the choice on purpose.

> ***The choice between Human-in-the-Loop and AI-in-the-Loop is the biggest design decision in your AI programme. Make it deliberately, or your vendor will make it for you.***

If you want to see what that map looks like in your own organisation, book a Workforce X-Ray pilot in a single function. It is the map that makes this decision possible.

This is the seventh article in a series drawn from the Bloor Research white paper *FusionWork™: Defining a New Category of Work*. The first asked why approximately 90% of job descriptions no longer describe the job. The second introduced the Thirds Model and the work nobody writes down.

ARTICLE REPRINT

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