

The Transformation Brief

Issue 18 · Friday 7 August 2026

Craig Horton

The Transformation Brief

Good morning, on Sunday the next provisions of the EU AI Act came into force, including transparency obligations for certain AI-generated content and the Commission's enforcement powers for general-purpose AI providers, with the high-risk rules covering employment, credit and other sensitive decisions now applying from 2 December 2027.

The extension gives leadership teams more time to prepare, although the agents already screening candidates, handling customer information and moving work between systems continue to operate today, which brings me back to something I have been asking for a while, can you count your agents?

The lead signal

Can you count your agents?

[SAP LeanIX published research this week](#) covering 228 enterprise architects and IT decision-makers, with 44% of companies already deploying AI agents, 40% experimenting with them and another 14% planning to, yet only 17% have visibility of agent performance or conformance, while 48% have no clear roles or responsibilities for managing them.

SAP calls this agent sprawl, which happens when agents are created, deployed or connected across systems faster than the organisation can inventory them, assign ownership, control permissions, monitor their behaviour or retire them when they no longer fit.

We have seen this pattern before with shadow IT, except an agent can carry credentials, reach across systems, call other tools and take actions at a speed most approval processes were never designed to govern.

For every agent operating today, I would want its register to answer five questions:

- Which systems does it touch, and what data can it reach?
- Which decisions can it take without a human?
- Whose budget does it sit on, and who approved it?
- What evidence does it leave behind for an auditor?
- Who has the authority to restrict it or switch it off?

Before the next agent goes live

An agent register should answer five questions

A board cannot govern an agent it cannot trace to a system, a decision and a named owner.

1

Systems and data

Which systems does it touch, and what data can it reach?

2

Decision authority

Which decisions can it take without a human?

3

Budget and approval

Whose budget does it sit on, and who approved it?

4

Evidence

What does it leave behind for an auditor?

5

Switch-off authority

Who can restrict or stop it when the work goes wrong?

The register entry

One agent, one record, one named owner with the authority to act.

Written in a personal capacity from public sources. Views are my own.

brief.craighortonadvisory.com

The five fields I would want completed before the next agent enters a live workflow.

Most organisations can name the platform and explain the task, then the answers become less clear when ownership, decision authority and evidence are discussed, especially when the agent crosses from one function into another and the outcome no longer belongs neatly to IT, HR, Legal or the business.

The delayed high-risk deadline has given organisations time to complete this work, and I would use that time to ask for the inventory, permissions, decision rights and named owners before the agent estate becomes too large for anyone to understand properly.

On the radar

1 Sandoz is preparing its managers for nearly 2,000 agents

Sandoz is implementing nearly 2,000 AI agents across the business, while Chief People Officer Tripti Jha is building a common organisational language around AI, role-based certification and an AI Coach that helps managers prepare for important conversations.

Nearly 2,000 agents will change how work moves through the company, and its managers will have a large part in whether people understand the change, trust the work and know where their judgement still has authority.

2 One organisation has put real money behind change

[The Medical University of Lusatia](#), a new organisation being built with EUR 3.6 billion in funding, has allocated 4.3% of its total budget to change management, making it the largest single line in its SAP contract.

I found this interesting because the budget covers the work around the system, including communication, learning, local ownership and support after go-live, with adoption funded as part of what the organisation is building from the beginning.

3 Policy is getting closer to the work

[Red Hat launched asago this week](#), an open-source project designed to translate AI policies into risk tests, mitigations, runtime controls and an audit trail connecting individual requirements with evidence from the operation.

It is early and we still need to see how it performs in practice, although the direction makes sense to me, because organisations need to show how each policy becomes a control, how that control is tested and who owns the decision when an agent acts.

The question I would take into the room

Before your next agent goes live, can you produce its register entry, show the decisions it can take without a human, name whose budget it sits on, explain the evidence it leaves behind and identify the person with the authority to switch it off?

Once again, I appreciate your interest and time in reading this brief, and as always, I look forward to your feedback and comments. Enjoy your weekend, Craig

Written in a personal capacity using public sources. The views are my own and do not represent Pax8. No confidential customer, partner or company information has been used.

The Transformation Brief · brief.craighortonadvisory.com · [Archive](#)