The Transformation Brief

Issue 04
24 April 2026

Lead Signal

Manual oversight has reached its breaking point.

From supervising tasks to governing systems.

Craig Horton Advisory

brief.craighortonadvisory.com

Lead Signal

Our reliance on manual oversight has reached its breaking point, as agentic workflows scale, we must move from supervising tasks to governing systems.

Deloitte's 2026 State of AI finds **85% of enterprises** are customising agents in production, while only **21%** have a mature model for governing them. The EU AI Act high-risk regime activates on **2 August 2026**, 102 days from today. Anthropic's Project Glasswing disclosed this month that Claude Mythos Preview autonomously surfaced thousands of zero-day vulnerabilities before public release, the clearest evidence yet that human review cannot keep pace with agent discovery speed.

From	To
Supervising tasks	Governing systems

The common thread, the agents have scaled faster than the governance designed to contain them, and the enforcement deadline is no longer theoretical. Every enterprise running agents in production without a documented governance layer is carrying a liability that compounds weekly until August.

Earlier this week I put aibvf-mcp into the MCP registry at v0.2.0 for exactly this reason, a pre-flight check inside any Claude agent before it recommends an AI initiative, returning Accelerate, Fix, or Stop against published benchmarks. Repo at github.com/Bahamas1717/ai-bvf.

Sources Deloitte, "State of AI 2026, Governance Gap Widens", April 2026. · McKinsey, "State of AI Trust 2026", April 2026. · Gartner, "AI Governance Platform Market", 7 April 2026.

On the Radar

Five signals from the week.

Each signal is important enough to be on the C-Suite agenda this quarter.

Anthropic · April 2026

$0.08

per session hour

Anthropic restructured Managed Agents this month, flat-rate seat pricing is gone, replaced by monthly commitments sized to estimated token use. The architecture separates reasoning from execution, a decoupled Brain and Hands with crash recovery from a persistent session log, and cached prompts get a 90 percent discount, 0.50 dollars per million tokens against 5.00. AI spend has moved from a headcount line to a capacity planning line, and every Q1 contract on flat-rate pricing is in Q2 scope for review.

Salesforce · April 2026

360

new release, four modules

Salesforce released Agentforce 360 this month, four components, Builder for agent development, Script for behaviour control, Voice for customer conversations, and Intelligent Context for grounding agents in unstructured data. Read together, the workflow reconstruction toolkit productised, Salesforce has moved from selling an AI module on top of existing workflows to selling the kit that builds the next generation of workflow inside its own stack.

McKinsey · April 2026

97 / 36

exploring vs governing

McKinsey's State of AI Trust 2026 finds 97 percent of enterprises running agent experiments while 36 percent have centralised governance, and the 61-point gap is the void that Project Glasswing exposed this month, with Anthropic's Claude Mythos Preview autonomously surfacing thousands of zero-day vulnerabilities before public release. Human review does not scale at Gen 3 discovery speed, and the regulator arrives in August whether the governance model is ready or not.

EY · April 2026

130,000

auditors across 150+ countries

EY has embedded a multi-agent framework into EY Canvas, 130,000 auditors across 150+ countries and 160,000 engagements now processing 1.4 trillion lines of journal data per year. This is the agentic consulting era made concrete, and the governance question it surfaces is live, who is accountable when an AI recommendation in a statutory audit is wrong. When the firm grading your controls has shipped agents before you have, the pressure to stand up your own agentic governance model compounds from Monday onwards.

BCG AI Radar · April 2026

1,7%

of revenue invested, two-thirds not scaling

BCG AI Radar 2026 reports corporate AI investment has doubled to 1,7 percent of revenue, CEO-led and board-visible, and two-thirds of enterprises have still not begun scaling AI across the enterprise. Doubling investment while failing to scale is escalation of commitment, more money into the same structural blockers, and the operating model is where that scaling gap actually lives.

Sources Anthropic Managed Agents pricing announcement, April 2026. · Salesforce, Agentforce What's New, accessed April 2026. · McKinsey, State of AI Trust 2026, April 2026. · EY newsroom, enterprise-scale agentic AI audit, April 2026. · BCG AI Radar 2026, CEO investment strategy, April 2026.

Against the Grain

> The consensus line this week is that Agent 365 and its peers deliver "governance at the SKU level". That is a vendor narrative, not a governance model. A control plane you can buy is the tooling layer, the governance work sits above it, in decision rights, accountability, and the operating model that decides when an agent acts and when a human intervenes. Treating the SKU as the answer is how organisations arrive at the EU AI Act deadline with a licence and no log.

From → To

The week's signals converge on one structural move. The organisations making progress are moving from supervising tasks to governing systems, from treating AI as an add-on to rebuilding the workflow around it, and from investing in models to redesigning the operating model that runs them. EY has shipped agents to 130.000 auditors, BCG has measured CEO investment doubling while scaling stalls, and McKinsey has put numbers on the 61-point gap between exploring and governing. Each signal is the same underlying transition, reported from a different corner of the market. The next six months are defined by who can build the governance layer at the speed the agents are already running.

> **The Reframing Question**
>
> What is the first agent in your organisation you would not let a regulator read the log of today?

If you cannot read the log, you cannot govern the agent.

Three governance questions for this week

- If an agent made a decision that cost you money this week, can you name who is accountable?
- If a regulator asked for an agent activity log tomorrow, could you produce one?
- Who owns each agent's outcomes by name, not by team?

What's Next

Next Friday, the EU Digital Omnibus trilogue on 28 April. We cover what the proposed AI Act amendments mean for high-risk classification and whether the August enforcement line holds. Microsoft 365 E7 also ships on 1 May, and we put Agent 365 to the governance test.

Craig Horton Advisory
AI Transformation, Enterprise Strategy, Change Management · linkedin.com/in/craig-horton-ai · brief.craighortonadvisory.com