

Craig Horton Advisory

The Transformation Brief

Issue 14 · Friday 10 July 2026 · brief.craighortonadvisory.com

This week

The Operating Model Change

4,800 roles are going. The operating model is the real story.

By Craig Horton, AI Transformation Advisor

The Operating Model Change

4,800 roles are going. The operating model is the real story.

This week Microsoft announced it is eliminating around 2.1% of its global workforce, about 4,800 people, as part of its company transformation.

Amy Coleman, Microsoft's Chief People Officer, said those roles are not being replaced by AI, and I believe that statement, while also believing the operating model is changing underneath it.

AI changes the volume of coordination, the decisions a manager sees first, the work that moves between teams, and the judgement leaders are expected to make when an agent is now part of the workflow.

A question came to mind: which work has genuinely changed, what decisions have moved, who owns the outcome once an agent is involved, and where does human judgement still have authority to intervene, challenge or stop the work?

This is where many AI programmes are heading, because now the operating model is starting to catch up.

From my office this week


Using the same agentic HR initiative from issue 13, I uploaded through the AI BVF, this time with the regulatory lens switched on.

Result, came back as Fix, leaning towards Stop, held by governance risk, because nothing should move until this was fixed.

The plan it returned was a pre deployment sequence:

- Classify the system under the EU AI Act
- Map every decision that affects a person, including where GDPR Article 22 may apply
- Design the human-in-the-loop role with genuine override authority
- Complete the data lineage and model-risk documentation
- Run the DPIA, remediate the findings, then re-score

It also calculates your waiting time, about EUR 5m to EUR 9m of organisational drag over a 10 week window, remember, that cost still applies even with your deadline pushed out.



The AI BVF regulatory readout from this run, the verdict, the four pillars against their targets, the binding constraint, the sequence, the cost of waiting and the re-score gate. The platform is open on the MCP registry as aibvf-mcp version 0.12.1, at github.com/Craig-Horton/ai-bvf.

On the radar

- Trust in AI agents has reached the standards agenda, with the International Telecommunication Union launching work this week on agent identity, accountability and meaningful human control, because agents are becoming operating actors and governance has to move into the operating model and design.

Reuters, 9 July 2026 · [ITU newsroom, itu.int](https://www.itu.int/newsroom/)

- The AI estate is spreading faster than oversight, with DigiCert reporting AI-related security incidents or vulnerabilities in 78% of organisations surveyed, alongside incomplete inventory and traceability, and a board cannot govern agents it cannot see.

DigiCert, 7 July 2026

- Infrastructure is now part of the value case, with Google Cloud reporting that 83% of organisations need material infrastructure change to capture the agentic opportunity, because data access, integration, compute cost and control architecture decide what reaches scale.

Google Cloud, 2026 State of infrastructure in the agentic AI era, 8 July 2026

Against the grain

"We are not replacing people with AI" may be true at the level of a role, but I am seeing enterprises change the work around the role: who prepares the decision, who reviews it, what evidence remains, and who carries accountability when the output is wrong.

That is why the redesign has to come before the project!

My reframe question

Before you approve a headcount action attached to an AI programme, can you show where the workflow changed, the decisions that moved, and the named person who owns the outcome?

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

Sources

- Microsoft, the latest in our company transformation, 6 July 2026 (blogs.microsoft.com)
- ITU, Focus Group on Trust and Identity for Humans and Agentic AI, 9 July 2026 (itu.int)
- DigiCert, AI security research, 7 July 2026 (digicert.com)
- Google Cloud, 2026 State of infrastructure in the agentic AI era, 8 July 2026 (cloud.google.com)
- AI BVF regulatory diagnostic, run 9 July 2026, aibvf-mcp version 0.12.1, github.com/Craig-Horton/ai-bvf