

Lead signal

The management layer drives twice the AI impact of the tooling.

Microsoft measured the readiness multiplier, a 67 to 32 split.

Lead signal

Microsoft measured what I have argued structurally for two years, the management layer drives roughly twice the AI impact of tooling, a 67 to 32 split.

The 2026 Work Trend Index measured the mechanism. McKinsey's 2025 State of AI logged the proof, **88 percent adoption** beside fewer than 10 percent of organisations running agents at scale. BCG puts **70 percent** of transformation value in people and process, Prosci puts failure six times higher when change-management spend drops below 10 percent of programme budget.

This is the Transformation Readiness Multiplier as a single statistic. An executive team funding only the 32 percent side returns a structurally smaller share on the same spend.

From	To
The tooling, 32 percent	The readiness, 67 percent

From my office, the agentic workmate layer went live on the Workflow Coroner this week, the senior-leader on-ramp that scores a workflow on the AI BVF and returns Stop, Fix, Accelerate, or Keep with a Monday plan named to the workflow. aibvf-mcp adoption held at **1,142 npm downloads** through the rolling thirty-day window, sustained signal not a one-off burst.

The procurement-flow fixture eliminates 11 of 14 steps for a 2.76M EUR annual saving on a 2.4bn EUR firm. Live at brief.craighortonadvisory.com/workflow-coroner.html, source at github.com/Bahamas1717/ai-bvf.

On the radar

Four signals from the week.

Each signal is on the executive agenda this quarter.

Stanford AI Index · April 2026

20%

drop in junior developer employment

Stanford economists show employment for software developers aged 22 to 25 fell nearly 20 percent since 2022. Agent success on real-world tasks moved from 20 percent in 2025 to 77.3 percent in 2026. The workforce story McAfee predicted in Issue 05 now has hard data attached.

OpenAI and Anthropic · May 2026

\$5.5bn

integration layer absorbed in seven days

OpenAI launched DeployCo on 11 May, 4 billion dollars from 19 firms led by TPG, plus Tomoro adding 150 forward-deployed engineers. Anthropic stood up a 1.5 billion dollar Wall Street joint venture led by Blackstone a week earlier. The reversibility cost when the vendor's services arm is fired in eighteen months is the new line on every brief.

IBM and ServiceNow · 5 May 2026

1 in 5

companies ready to govern autonomous agents

IBM at Think 2026 shipped a governance-first operating model blueprint on 5 May. ServiceNow at Knowledge 2026 bundled its AI Control Tower across every product. Deloitte 2026 finds only one in five companies has a mature model for governing autonomous AI agents, three control planes shipped against that gap.

Council of the EU · 7 May 2026

18 mo

the high-risk cliff edge moves later

The Digital Omnibus reached provisional agreement on 7 May. High-risk deadlines now sit at 2 December 2027, embedded AI in regulated products at 2 August 2028. The cliff edge slipped, the substance held, the risk is wasted compliance theatre. Executive teams should still design to the high-risk standard.

Against the grain

The 67 to 32 split is real, but the question it asks is incomplete. Microsoft measured organisational readiness against the existing workflow, the workflow agents are about to redesign sits outside that frame. Executives funding the 67 percent side without asking whether the workflow should exist are buying readiness for work the agent layer will eliminate within twelve months.

Readiness amplifies whatever sits underneath. Workflows worth keeping return more on the spend. Score the workflow first, then retire those that don't.

Perspective

Microsoft measured what BCG and Prosci have argued structurally for two years, organisational readiness drives roughly twice the AI impact of tooling. The buyer's question this quarter sits upstream of platform selection, the readiness that carries two thirds of the return, the partner that owns the deployment hours, the operating model around the model. Any executive team still treating AI as a procurement question will spend the budget twice.

The reframing question

If two thirds of your AI return depends on conditions you have not funded, what is your pilot actually measuring?

Two tests this week to find the answer.

Run these against the programme budget

Map the 67 to 32 against your current AI programme budget. If the spend split inverts the value split, the maths is broken.

Calculate change-management spend as a percentage of programme spend. Below 10 percent is Prosci's red zone, the multiplier sits underneath.

What's next

Issue 07 lands Friday 22 May 2026. The Chief AI Officer, where it shipped without the mandate, and the five-question diagnostic that separates a mandate from a side project.