

CHANGEMAKERS | AI LEADERS ROUNDTABLE

Key insights

From experimentation to operation: making AI work in practice

Anonymised insights from a Changemakers roundtable. Names and organisations have been removed so the learnings can be shared openly.

KEY INSIGHTS

From the messy middle to AI that works

Most marketing teams have AI switched on. Copilot (or an equivalent) is rolled out, a handful of agents exist, and someone is quietly wondering whether any of it is actually solving anything. That is the messy middle: enough activity to look busy, not enough structure to know if it is working.

Further along in the AI journey, a different picture emerges: teams with token usage they can actually measure, an AI source of truth they can trust, and a KPI story that survives contact with the board. Getting from one state to the other does not need a bigger AI budget. It needs a plan.

This report covers four themes: knowing when to stop and audit rather than build another agent, a simple framework for choosing where to start, the governance work that makes AI answers trustworthy, and the measurement question every board keeps asking.

IN THIS REPORT

01 **Audit before you automate**

02 **Build purposeful agents**

03 **Trust is the real bottleneck**

04 **Measure AI performance for the board**

THEME 1

Audit before you automate

The instinct across most marketing functions right now is to keep building. Tone of voice agents, market research agents, content agents: the tools multiply faster than anyone can judge whether they are earning their keep. The more useful instinct is to stop, audit what is actually broken, and only then decide whether AI is the fix.

Recognise FOMO for what it is

Speed pressure is real, but it is driving teams to build volume rather than value. The mandate to "do something" with AI is not the same as a strategy for what to do with it.

> "We are in this total AI crazy FOMO, everyone wants 5,000 agents for this and that."
>
> **Director, design and engineering software company**

Check whether automation solves it first

A chunk of what gets labelled an "AI problem" is actually an automation problem, and a much cheaper one to fix. The better approach is to audit what is actually causing friction, then work out whether the right answer is AI, automation or neither.

> "It wasn't really solving problems. We were just using AI for AI's sake."
>
> **CMO, HR and payroll software company**

Replace deployment thinking with problem thinking

Maturity is not measured by how many tools are live. It is measured by how clearly a team can explain why each one exists.

> "Think about the problem that you're trying to solve, and how AI will solve it, not where you can deploy AI."
>
> **Marketing leader, financial services organisation**



THEME 2

Build purposeful agents

Once the audit question is settled, the next challenge is prioritisation: with limited time and a long list of possible use cases, what goes first? One example is a framework known as MOTIVE, which breaks the decision into six questions: manual effort, ownership, tedium, impact, variability and existing data. Score a potential process against all six, and the highest-scoring, lowest-friction option becomes the obvious starting point.

Start the race before trying to win it

Analysis paralysis is as common a failure as reckless deployment. A small, well-scoped pilot beats a perfect plan that never ships.

> "You're not going to get to the finish line unless you've started the race."
>
> **AI implementation practitioner, marketing consultancy**

Pick the process people already dislike doing

One example is a sales team's cold outreach process, which scored highly on every measure: high manual effort, clear ownership, genuine tedium, high potential impact, low variability, and usable existing data in the form of case studies and account histories. That combination made it the obvious low-risk, high-value starting point. It was not the most ambitious idea on the list, but it was the easiest to build and the most likely to stick.

Measure adoption, not ambition

The resulting tool was built and piloted with a small group over eight weeks. A year on, it is still in daily use, with no requests to change or switch it off. The lesson generalises: success for a first AI project looks like sustained use and positive feedback from the people doing the work, not a headline ROI figure.



THEME 3

Trust is the real bottleneck

Access to AI tools is rarely the constraint any more. Trust in what they produce is. Once an AI-generated answer (a stat, a claim about revenue, a fact about a client) gets copied into an email or a deck, an untruth stops being a minor error and starts being a liability.

Don’t mistake plausible for true

Every team with broad AI access is exposed to this risk, whether or not they have noticed it yet. The fix is not restricting access. It is building a layer of truth that AI can be pointed at.

> “Those plausible answers are getting us in trouble, potentially, down the road.”
>
> **AI implementation practitioner, marketing consultancy**

Build a source of truth, not just a skill

Building truth into individual AI skills tends to drift over time. The more durable fix is a governance layer that ingests everything (documents, transcripts, case studies, video), de-duplicates conflicting figures, and lets any answer be traced back to its original source. Once that layer exists, new use cases such as events, tone of voice and client-facing tools can be built on top of it with far less risk.

Use guardrails to enable adoption, not slow it

Rolling out guardrails first, before wider adoption, is what makes people comfortable using AI in the first place. Far from being a barrier, governance built into the tools gives teams the confidence to use them.

> “It’s built in, so everyone felt they could use it.”
>
> **Marketing leader, financial analytics software company**



THEME 4

Measure AI performance for the board

Every AI investment eventually gets a KPI question from the board. The teams making progress have stopped trying to force AI into a single time-saving or cost-saving metric, and have started building a broader measurement story instead.

Know what is actually moving the metric

A single proxy metric, chosen for simplicity, can end up telling the wrong story. One example is a board-level AI metric, tracked closely as adoption picked up, that moved in the wrong direction entirely. AI was helping, but there was no reliable way to isolate its effect on a genuinely messy dataset.

Combine pipeline, ramp time and cost savings

One example of a broader measurement approach centres on pipeline per marketing employee, benchmarked against industry ranges from efficient through to elite performance. It sits alongside SDR ramp time and infrastructure cost savings identified through a full tech stack review. That review, using AI to assess what the stack could lose and where the savings should be reinvested, became the basis for reshaping the marketing organisation.

Redistribute skills before you cut headcount

Not every AI-driven change is a restructure. One example is a centre of excellence handling research, account planning and content for an ABM programme, which moved that work into an orchestration layer without reducing the team. The people who had been doing manual research and analysis were repositioned to manage, validate and refine the AI agents. Headcount stayed flat, but what people spent their time on changed completely.



Treat token cost as a planning input, not an afterthought

As more organisations move to enterprise-tier AI licensing, token consumption becomes a real budget line, not just a technical detail. In one example, early data suggested typical usage was a small fraction of the allowance provided per employee. The more important shift was building tokens into the ROI model from the start, rather than treating them as a cost to worry about later.

SUMMARY

Key takeaways

01 **Audit first.** Most AI problems are audit problems in disguise, and a chunk of what looks like an AI gap is actually an automation gap.

02 **Score before you build.** A simple framework across manual effort, ownership, tedium, impact, variability and existing data turns a long wish list into an obvious starting point.

03 **Treat governance as the unlock, not the brake.** A traceable source of truth is what lets teams scale AI use with confidence, not what slows them down.

04 **Pick metrics that survive scrutiny, and know what is driving them.** Pipeline per employee, ramp time and infrastructure savings tell a more honest story than a single proxy metric ever will. Restructuring and redistributing skills are both legitimate ways to act on that story.

None of this is really about technology. The tools will keep getting faster and more capable regardless of what any one team does. What separates teams making real progress from those stuck in the messy middle is discipline: the willingness to audit before building, to start small and prove value before scaling, and to be honest about what a metric is actually measuring.

That is a mindset shift as much as a technical one, and it is available to any team, regardless of how many licences they have already bought.



About Changemakers

Peer-led roundtables for marketing leaders navigating change. Honest conversations, practical ideas and anonymised insights you can act on.

About Modern

The B2B marketing consultancy for marketing leaders in complex businesses going through change.

We bring focus to what matters next, and the evidence to defend your decisions.

Accountable for the outcome, not just the output. We rebuild the fundamentals, from targeting to measurement, and leave your team able to run them.

Let’s connect

W: modernb2b.co
T: +44 (0)117 332 6700

Nicola Ray
nicola.ray@modernb2b.co
Connect on LinkedIn