

CHANGEMAKERS | ABM LEADERS ROUNDTABLE

Key insights

AI in ABM: from experimentation to agentic scale, with Kyndryl

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

KEY INSIGHTS

From piloting to scaling: where ABM leaders are with AI

AI has moved off the experimentation agenda and into the operational core of B2B marketing. For ABM this matters more than most: the precision ABM demands, from deep account intelligence to personalised content and timely activation, maps directly onto the tasks AI can now accelerate at scale. The challenge is not whether to adopt it, but how to do so in a way field teams will actually trust and use.

Most ABM leaders are still in pilot mode, experimenting with individual tools and use cases. A smaller number are scaling. One example of going further is Kyndryl, which is building a fully orchestrated, multi-agent AI system in-house with a team that started with no technical expertise. Its journey provides the practical backbone for the themes that follow, and shows where the real barriers between piloting and scale tend to sit.

This report covers four themes: getting started with the right use cases, building agents you can trust, driving team adoption, and planning for continuous optimisation and governance.

IN THIS REPORT

01 **Get started with the right use cases**

02 **Build agents you can trust**

03 **Drive adoption through power users and community**

04 **Plan for continuous optimisation and governance**

THEME 1

Get started with the right use cases

AI adoption stalls when teams reach for technology before understanding the work it is supposed to improve. The programmes that scale start with the workflow, prove one use case at a time, and only then bring the pieces together.

Map the work before building anything

The more productive starting point is understanding how ABM work is currently done: where time is spent, where effort is manual, and where output is slow or inconsistent. Field marketers typically carry a heavy research and profiling burden, and understanding a large enterprise account well enough to run a strategic programme can take days. That is exactly the kind of concentrated, repeatable task where AI makes an immediate difference.

Mapping the work does two things. It produces a prioritised list of use cases where AI can genuinely reduce effort, and it sets a quality benchmark so teams can judge whether any given output is good enough to use.

> "We looked at each of these steps and understood from our field marketers what they spend their time doing. Then we started breaking that down at a granular level. If I'm looking at account research, how do we currently go about doing it? We mapped the as-is state, and we realised there was a lot of manual effort involved, so how do we use AI to augment that?"
>
> **VP of ABM, Kyndryl**

Solve one use case at a time, then build on it

Once use cases are defined, the approach that works is iterative and narrow. Pick one, build a workflow or agent to solve it, test with real users on live work, gather feedback and refine. Only once that use case is working reliably should the next one be tackled. Account research is a good place to begin: refined over multiple rounds of testing with the people who rely on it, the time saving is significant.

4 to 5 days → 15 minutes
account research time with a validated AI agent

> "We started off rudimentary: a basic research of an account. Then we took it to the market, worked on live deals, live accounts, with the deal teams and field marketers, to get constant feedback. Is the information enough? Is the information correct? Each of these individual bots went through dozens of iterations to reach a point where the outcome was approved and vetted by the account partner who knows that account well."
>
> **Head of Global ABM Centre of Excellence, Kyndryl**

Move from individual bots to one orchestrated agent

As individual use cases are validated, a new problem emerges: ten separate tools for ten separate tasks create their own friction. The natural next step is a single orchestrating agent that understands a query, calls the right capability and returns a structured output, all through one conversational interface. That shift from a collection of bots to one orchestrated agent is what moves an AI programme from useful experiment to operational infrastructure.



THEME 2

Build agents you can trust

Early agent outputs are unreliable, and compliance teams are right to be cautious. Trust is built through sustained work with field experts, guardrails designed in from the start, and data added in stages as reliability is proven.

Train agents with the experts who know what good looks like

What turns unreliable output around is not better prompting in isolation, but sustained collaboration with the field experts who know what good looks like. Often it means dozens of rounds of refinement, testing live on real accounts with deal teams, and only signing off outputs when the account partner confirms the intelligence is valid. In this example, Microsoft Copilot was the main engine used.

Training on examples of what good looks like is particularly effective for structured outputs. One example is a bot trained on 20,000 historical RFPs, mapping internal and external examples to define a quality benchmark before the agent produced any live output.

1 week → 50 minutes

proposal evaluation cycle with a trained RFP agent

> "We worked with actual people in the field who know those things very well. They act as a barometer to train the agent on what best or good quality looks like. We do the iterations again and again, keep telling it: no, this is not what I want, I want it like this. Then you get the next version, and work with your field experts until they say: yes, this looks like what I can use."
>
> **Head of Global ABM Centre of Excellence, Kyndryl**

Design compliance in early, as a gate, not an afterthought

The most common risk areas are access to first-party contact data, exposure of personally identifiable information (PII), and publicly available information used in ways that may not meet data protection requirements. Just because data is publicly available does not mean there is legal permission to act on it, and that industry assumption is likely to attract regulatory scrutiny.

The teams moving fastest through compliance review do the work early: documenting guardrails, limiting access to sensitive internal systems, and walking legal and CIO teams through the detail before seeking production approval. Where there is a top-down commitment to AI transformation, the commercial case rarely needs making. What earns production approval is the due diligence.

> "We consciously made an effort: contact databases that we own, agents cannot access that. It is off-limits to them. Persona-related modules look at a person's work commitments and public statements. No email addresses, no phone numbers, whether verified or not."
>
> **Head of Global ABM Centre of Excellence, Kyndryl**

> "We have never been asked to prove it. The mindset is there from top down. We have to be an AI-native marketing organisation. What we had to justify was to the CIO, in terms of compliance. We had to prove that we had done the due diligence: negligible hallucination, no data breach threat, no PII being exposed."
>
> **VP of ABM, Kyndryl**

Start with public data, then layer in internal data once trust is proven

The first version of an agent does not need to connect to everything. Build on external, public web sources first and use that deployment to prove reliability before wiring in CRM, intent platforms and customer data. Whitelisting known, dependable sources reduces hallucination significantly at this stage.

When intent data is added, combining sources works better than relying on one signal. First-party website activity, seniority scoring of individuals showing intent, and third-party signals layered in a weighted scoring model produce far more reliable results than third-party intent data alone.

> "Getting the output in a validated, negligible hallucination state from external sources is the actual tough part. Once we solved that, we started adding the layers of internal data, which will enrich all the outputs."
>
> **Head of Global ABM Centre of Excellence, Kyndryl**



THEME 3

Drive adoption through power users and community

Deploying AI agents is the easy part. Getting a distributed team to use them consistently and effectively is where most programmes run into difficulty.

Use power users as your fastest route to adoption

Rolling out to a limited beta cohort and tracking usage from the first two weeks reveals clear gradations: power users, moderate users, and those who have not engaged at all. Each needs a different approach, and power users activated as internal advocates drive adoption faster than any formal training programme.

> "Someone who is a power user, we use them as our spokesperson across the community. People who are not very comfortable, we go on a much easier scale: let us handhold you into this. Start with three or four use cases. Are you comfortable? Now try something else. And it is a continuous process."
>
> **Head of Global ABM Centre of Excellence, Kyndryl**

Curate the agent estate before it sprawls

Adoption can also swing the other way, when users start to create their own bots and workflows in silos. Agent proliferation, sometimes called "agent slop", emerges as individuals build without visibility of what already exists. Active curation, monitoring and governance, along with training existing agents rather than building new ones from scratch, keeps the tool estate manageable and quality consistent.

Build a community, not just a training plan

Peer advocacy is a more powerful enablement mechanism than formal training. An internal community where people share wins and challenges tends to drive more sustained engagement than top-down instruction, although instruction still helps.

> "Every week we have a community call for the field marketers doing ABM. That will consist of our team sharing the latest capabilities, or it will be one of those field marketers talking about a deal they just won, or a challenge they have. Getting that community working is a really good way of getting enablement out there. It is not people telling them what to do, it is enabling people to help themselves and each other."
>
> **VP of ABM, Kyndryl**

Share tools with forward-thinking commercial teams

AI in ABM also challenges the traditional boundary between marketing and sales tooling. Where commercial teams are forward-thinking, sharing tools with them rather than withholding them helps identify further use cases and drives adoption.

THEME 4

Plan for continuous optimisation and governance

AI deployments are never done. New use cases emerge from unexpected users, new requirements surface once adoption broadens, and the technology keeps developing. Planning for that from the outset changes how teams resource and govern their programmes.

Plan for continuous iteration from the start

Treating any release as a finished product sets the programme up to stall. The teams making progress plan and resource for ongoing optimisation, expecting each release to surface new use cases and new failure points.

> "When it comes to AI, we will always be in the optimisation stage. You reach a point where you feel it is ready, it starts being used, and then you get new use cases, new fault lines showing up, which you need to correct and make better. I do not see an end state."
>
> **Head of Global ABM Centre of Excellence, Kyndryl**

Build governance into the operating rhythm

In practice, governance means access controls that respect individual permission levels, active curation of the agent estate to prevent proliferation, and regular feedback loops with field teams. These belong in the normal operating rhythm, not in a list of one-off project tasks.

Start small, and expect it to take time

Going from scratch to a fully orchestrated agent takes time, typically 12 to 18 months. The key is to start, and to start small: solve for one use case and build out from there.



SUMMARY

Key takeaways

01 **Start with the work, not the technology.** Map how ABM tasks are currently done, identify where manual effort is heaviest, and let use cases emerge from that. Account research is one of the most impactful starting points: with a validated agent, a process that takes days can be completed in under an hour.

02 **Build one use case at a time and prove it before scaling.** Research and preparation time can fall by up to 90%, proposal evaluation cycles can drop from a week to under an hour, and content customisation time can fall by around 80%. None of that happens without first proving reliability at the individual use case level.

03 **Compliance is the bottleneck, not the business case.** Where leadership has committed to AI transformation, the commercial argument is already made. The work is getting agents through CIO and compliance approval: design guardrails early, keep first-party contact data off-limits in the first iteration, and walk stakeholders through the detail before seeking production sign-off.

04 **Power users and community beat formal training.** Segment beta users by engagement level from day one, use power users as advocates, and build a peer community where wins and challenges are shared.

05 **Optimisation is permanent, so resource it accordingly.** Every release will surface new use cases and new failure points. The teams making the most progress treat iteration as an operating model, not a project phase.

The skills that matter are less about knowing which platforms to buy and more about curiosity, a willingness to experiment, and the discipline to involve field teams at every stage.

Marketers building the most useful AI infrastructure are not technical specialists. They are people who understand the commercial context deeply enough to know what good output looks like, and are persistent enough to iterate until they get there.



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