




Scaling Agentic AI

Why better tools aren't enough, and what moves organisations forward





Scaling Agentic AI

Executive summary

Leaders want agentic AI: over half the practitioners we surveyed feel the push from the top. Today's AI makes people faster. Agentic AI can reshape how the whole organisation works. The companies pulling ahead get five things right, the '5 Ps' of permission, problem, personas, proximity, and proof, and give agents as much freedom as a mistake can afford. Start with one problem, one persona and permission to try.

Somewhere in many organisations right now, someone senior is asking for agentic AI: a formal mandate from the top, or an increasingly confident line in a leadership meeting. Either way, the ask rarely comes with much detail attached, and it usually assumes a great deal has already been worked out that, in practice, hasn't.^[1]

And yet the pressure to adopt AI, especially agentic AI, in M&A and integration is intense: 86% of corporate and private equity leaders say generative AI is now built into how they run deals,^[2] as acquirers look to move faster, cut cost, and keep pace with competitors already ahead of them.

We asked a small group of M&A and integration practitioners where they stood. More than half described strong or moderate encouragement from their leadership to adopt AI. Not one practitioner described their own organisation as running agentic AI autonomously today. That gap is the starting point for this paper, not the technology itself.

So here is a question worth sitting with before anything else: is fully autonomous AI what your organisation needs, or is it the version of the story that's easiest to tell? This paper argues for the second reading, cautiously, drawing on first-hand interviews at an organisation already taking deliberate early steps into agentic AI in M&A and integration and a pulse survey^[3] to set out where organisations stand today, what the more advanced ones are doing differently, and what a more typical organisation can do next.

Where are we, really, and what do we mean by agentic?

Before asking what's possible, it's worth being honest about where things stand, not compared with competitors, most of whom are dealing with the same questions. The better measure is how much of the technology's opportunity can be captured safely today.

Four categories describe most of what gets called AI today, and a lot of the confusion in this space comes from using the word 'agentic' for all of them at once (Figure 1):




Figure 1: Four categories of AI autonomy



- **Rules-based:** the same input always produces the same output, no judgement involved. Example: automatically flagging a contract when its renewal date falls inside the TSA exit window.
- **Assisted AI:** a person prompts the system and decides what to do with the answer. Example: drafting a first-pass summary of a data room folder.
- **Bounded agentic AI:** the system plans and acts across several steps, always inside a checkpoint a person controls. Example: drafting and checking a first-pass valuation analysis before a person signs it off.
- **Autonomous agentic AI:** the system completes a task end to end with nobody confirming before it acts. Example: negotiating and signing off a vendor contract with no review.

Our survey suggests most practitioners' own use today sits at Assisted AI, prompting a tool and deciding what to do with the answer. Our interviews suggest something more striking: even at an organisation built for speed and comfortable with risk, nobody has gone further than Bounded agentic AI in M&A or integration work,

not because nobody has tried, but because nobody has needed to. That organisation has developed due diligence and integration tools built around a series of specialist agents: one drafts, one challenges, one checks, with a person confirming at every handoff. Full autonomy stays out of reach on purpose. The team's own account is blunt: leave agents unsupervised and things will go badly wrong.^[3]

Moving an organisation from one category to the next takes five ingredients (the '5 Ps'): **Permission** to try things that might not work; **a specific, painful problem** to solve; **a persona**, someone that will benefit from the output; **proximity** to the people doing the work; and **proof** that the change will be effective.

Skip any one of those five and progress stalls, however good the model becomes. Moving towards full autonomy in particular needs a track record, a demonstrable, compelling business case or benefit, real performance evidence built up while a person was still checking, a documented decision rather than something that drifts (Figure 2).




Figure 2: The criteria to move up the scale



Why aren't we fully autonomous yet, and should we want to be?

That still leaves the question: with the pressure to adopt agentic AI this intense, how much of it does a typical organisation need? It's tempting to treat full autonomy as the natural finish line, the place every organisation is heading given enough time and confidence.

Our interviews suggest a different reading. Even at the organisation most willing to move fast and take risks, nobody could point to a fully autonomous agent running inside an M&A or integration process. What exists instead, everywhere, is a system that acts, with a person confirming before or after.

"It's not actively doing stuff that we can't control. All it's doing is surfacing huge amounts of information."

—Director, M&A Integration

That is a design choice. Two candid, first-hand mistakes shaped the current approach: a report that went out with an error nobody had caught, and a tool that quietly invented a person to make a number add up.

Both taught the same lesson: check everything AI produces, always.^[4]

The real question for most organisations isn't when will we reach full autonomy.

It's where, specifically, could a mistake be cheap enough that we'd even want to let go of the wheel? For much of integration work, people decisions, legal exposure, anything hard to reverse, that answer may reasonably be: not yet, and maybe not for a while.

So when is autonomous the goal?

Pulling together our survey, our interviews, and how this question is being answered elsewhere, autonomous is the right goal when three things are true at once:^[5]

- The task is high volume and bounded, not a judgement call dressed up as a process.
- Errors are cheap to catch and easy to reverse.
- There's already a track record, real performance evidence built up while a person was still checking.

Spotlight: Ocado

Outside M&A, one example shows what it looks like when all three conditions hold at once. Ocado's grid-based fulfilment centres run largely autonomously: swarms of robots move across the grid at up to 4 metres per second, picking and routing stock with no person directing each individual action. One 600,000 sq ft site processes around 65,000 orders a week with fewer than 600 staff. It works because the task is bounded, mistakes are cheap to fix, and the system has years of track record, not because the technology is more advanced than anywhere else.

Sources: FreightAmigo, 2026; Singularity Moments, 2026.



Where any one of those is missing, keeping a person in the loop isn't a temporary compromise. It's the appropriate design.^[6] Organisations working in regulated, high-stakes environments elsewhere already treat that as their permanent target state.

So what are other companies doing?

Agreeing what 'better' means. Ask a room of practitioners what they want AI to achieve and the answers vary:

- Faster deal execution
- Better target analysis
- More insightful due diligence
- Smoother experience for people joining on Day One

Each pulls an operating model in different directions, and few organisations have chosen between them. This is where the experience of one organisation in our research is instructive, because of the question they kept returning to: how do we make our processes and services more effective, not more efficient? Choosing between the two is a leadership decision, not a technology one.

What more advanced organisations do differently. Leadership backing for AI is common, our survey finds, but clarity about who owns the resulting decisions is not, and that gap is what separates breadth from depth:

a function can have one deep, carefully governed autonomous pilot while the rest of the business is still at Assisted AI, or the reverse.^[7]



Scaling agentic AI comes then down to five practices, all within reach of any organisation, without any exceptional technology or an exceptional appetite for risk:

- Budget for experimentation before proof, with formal governance following once it scales, not before.
- A sandbox for safe experimentation: real but harmless data, so people can try things without risk to a live deal.
- AI support specialists embedded inside each function, rather than one central AI department trying to serve everyone at once.
- A regular cadence for showing and sharing what's been built across functions, so techniques and tools get copied and improved rather than reinvented in every corner of the business.
- Cybersecurity and legal in the room from the start, not consulted only once a tool is already built.



Spotlight: The Access Group

During due diligence on a recent deal, a tool built to search across the full data room flagged a discrepancy buried in the target's own paperwork, a stated count of data centres that didn't match its supporting documents, live, in the meeting. That tool is now used on every acquisition Access makes. It's a small, concrete example of the pattern behind every point above: value showed up first in one narrow, well-bounded task, and only spread further once it had proven itself.

Done well, this shifts what a team can build, and how quickly an idea becomes something people use.

Now what should I do?

Individually, that starts with being honest about where your own function sits on the scale above and picking one bounded, low-stakes task to test the water. But making it stick needs five moves at the organisational level, and they build on each other.

- Define the why, before anything else. Agree, at leadership level, what 'better' means for your M&A and integration function.
- Empower employees to explore, test and trial in a safe environment. This is two things done together: a real sandbox, safe data, safe consequences; and a leadership message, repeated, that trying and failing at this stage is expected.

- Set up a distributed Centre of Excellence. Put a small number of specialists inside each function to make the tools usable, rather than one central team trying to serve everyone.
- Run a regular cadence for sharing what's been built. Bring functions together on a fixed schedule to show what they've tried, get challenged on it, and copy what's working elsewhere, so effort compounds instead of the same tool getting built several times over in different corners of the business.
- Involve cybersecurity and legal early, not at sign-off. Avoid a stalled project by bringing these functions into the room at the start of a new workflow, changing their role from gatekeeper to co-owner.





Where does this leave you?

You can start without a fixed view of where AI in M&A ends up. The organisations pulling ahead aren't the ones with the clearest five-year plan.

Give people time, tools and a budget to try things on live problems, and make clear that failed attempts are expected (Permission).

Start with one specific, painful problem and let each step inform the next (Problem).

Name the person who will use it before you build it (Persona).

Put the capability next to the people who feel the friction every day (Proximity).

Measure what gets faster, better, and extend autonomy as far as the cost of being wrong allows, and be ready to change direction as new benefit areas emerge (Proof).

References

- [1] Stanford HAI, 2026 AI Index Report.
- [2] Deloitte, 2025 M&A Trends Survey, cited in DealRoom, AI Due Diligence in M&A: Uses, Risks & Limits, 2026.
- [3] Oaklin interviews and pulse survey of M&A and integration practitioners, 2026.
- [4] McKinsey & Company, State of AI Trust in 2026: Shifting to the Agentic Era.
- [5] Gartner, 2026 Hype Cycle for Agentic AI; Gartner Predicts Over 40% of Agentic AI Projects Will Be Cancelled by End of 2027.
- [6] International AI Safety Report 2026.
- [7] McKinsey & Company, The State of AI: Global Survey, 2026.



Oaklin

Wild about
ideas, serious
about business

Want to learn more about M&A,
Integration and Inorganic Growth?



Scan the
QR Code

oaklin.com