

BEYOND THE PLAYBOOK

M&A Integration in the AI Era — The Playbook Gains a Chapter

What three decades of integration practice never had to account for — and why it matters now

Giovanni Leonardi

May 2026



M&A; Integration in the AI Era — The Playbook Gains a Chapter

Transformation · No. 438
May 2026

“The cost of deferral is no longer measured in weeks of human confusion. It is measured in thousands of automated decisions made under assumptions that may already be wrong.”

1. The War Room's Blind Spot

The integration war room looks the way it always does. Workstreams are tracked on the wall — IT systems, HR, finance, operations, commercial. Day-1 readiness items are green, amber, or red. The synergy model is loaded and the first hundred-day milestones are mapped. The playbook, refined over decades and dozens of transactions, is running as it should.

But nobody in the room has asked what happens to the target's fourteen production models when the data they were trained on changes ownership. Nobody has mapped the automation estate — the forty-odd workflows executing yesterday's processes at machine speed while the integration team debates tomorrow's operating model. And nobody has checked whether the acquiring company's AI governance framework can absorb, or even recognise, the target's approach to model risk.

This is not a failure of diligence. It is a gap in the playbook itself.

2. The Intelligence Layer Is Now Material

For most of the history of M&A; integration, the assets that mattered were tangible enough to inventory: systems, contracts, people, intellectual property, customer relationships. The integration playbook — Day-1 planning, systems consolidation, org design, synergy tracking — was built around these. It assumed, reasonably, that the things being integrated were largely inert between decisions: a system sits where it is until someone migrates it; a process continues as-is until someone redesigns it.

The reasonable response is that AI is just another technology stack — that models and automations are simply more systems for the IT workstream to absorb. If the playbook already handles ERP consolidation, CRM migration, and data-centre rationalisation, why should AI be any different?

It should be different because AI assets are not inert. The intelligence layer — the models, the training data, the automated decision-making, the agent-driven workflows — does not sit and wait. It executes, and in some configurations adapts, whether or not the integration team has made a decision about it. An ERP system that has not been migrated still runs yesterday's process, but it runs it the same way every time. A pricing algorithm that has not been recalibrated also runs yesterday's process — but it may be learning from today's data, optimising against assumptions that no longer hold, and compounding the drift with every batch. The distinction between a static system and an active model is not a technicality. It is the reason the playbook needs a new chapter.

Three things in particular are genuinely new.

First, data is no longer just a migration problem — it is a rights and provenance problem. In the traditional playbook, data migration is an IT workstream: move the records, reconcile the schemas, decommission the legacy system. But when the target has trained models on that data, the question is no longer just where the data lives. It is who has the right to train on it, whether consent and licensing terms survive the change of ownership, and whether the acquiring company's own models can legally learn from it. Training-data provenance — knowing what went into a model and on what terms — has become a due-diligence line item that most integration teams are not yet equipped to assess.

Second, the model estate is an asset class the playbook has no chapter for. Production models carry embedded assumptions about the data distribution they were trained on, the business context they were deployed in, and the guardrails that were placed around them. An integration that consolidates the underlying data without understanding the downstream model dependencies risks degrading or invalidating models the target's operations depend on. Model inventories — what exists, what it does, what it depends on, who governs it — should be as standard a part of integration planning as systems inventories became a generation ago.

Third, AI governance regimes collide. Every organisation that has deployed AI at any scale has made choices, whether deliberate or emergent, about how it governs that deployment: risk classification, human-in-the-loop requirements, monitoring, audit trails, bias testing, incident response. These choices reflect not only regulatory obligations — and the EU AI Act's compliance timeline is making several of them mandatory — but organisational culture and risk appetite. When two companies merge, their governance regimes meet. If neither side has documented its regime clearly enough to compare, the integration team inherits a silent liability: production models running under assumptions that no longer apply in the combined entity.

3. The Old Truth, Sharpened

There is a principle that experienced integration practitioners know well: deferred decisions are not cost-free. Every day that a decision about systems, people, or processes is delayed adds friction, uncertainty, and quiet cost. The organisation does not pause while the integration team deliberates; it continues to operate, and the gap between the current state and the intended future state widens.

AI sharpens this truth to a point. In the pre-AI world, deferred decisions accumulated cost at human speed — people worked around unclear structures, processes drifted, morale eroded. Ungoverned automation compounds at machine speed. Consider a pricing algorithm trained on the target's historical margin data. On Day 2 of the combined entity, the acquiring company adjusts the product portfolio — a routine early integration move. The algorithm, unaware of the change in commercial strategy, continues to optimise for the old margin structure. By the time anyone notices, it has repriced several thousand SKUs against assumptions that no longer apply. The cost of deferral is no longer measured in weeks of human confusion. It is measured in thousands of automated decisions made under assumptions that may already be wrong.

This is why the intelligence layer cannot be treated as a later-phase workstream. It must be part of Day-1 planning — not because AI is the most important thing in every integration, but because it is the thing most likely to keep moving while everything else waits.

4. An Intelligence-Layer Due-Diligence Checklist

The practical response is not to reinvent the integration playbook but to extend it. The discipline is the same — inventory, assess, decide, execute — but the objects of that discipline now include a category the playbook was never written for.

The following areas demand attention before Day 1, not after:

- Model inventory and dependency mapping. What production models exist? What do they do? What data do they depend on, and what business processes depend on them? This is the equivalent of the systems inventory that became standard practice in the early 2000s — except that models degrade silently when their dependencies change, where a system failure is usually loud.
- Training-data provenance and rights. On what data were the target's models trained? Under what terms was that data collected or licensed? Do those terms survive a change of corporate ownership? Can the acquiring company's own models legally train on the target's data, and vice versa? These questions sit at the intersection of legal, data governance, and AI — a junction most integration teams do not yet staff for.
- Automation estate mapping. What automated workflows — rule-based or model-driven — are currently executing? What processes do they encode? Which of those processes will change under the integration plan, and what happens to the automations when they do? The risk is not that an automation breaks. It is that it continues to work, faithfully executing a process that is no longer correct.
- AI governance regime comparison. How does each organisation classify AI risk? What human-in-the-loop requirements are in place? What monitoring, audit, and incident-response processes exist? Where do the two regimes conflict, and what does the combined entity's regulatory obligation require? Under the EU AI Act, this is not a matter of best practice. It is approaching a matter of law.
- Agent and autonomy audit. Where AI agents operate with any degree of autonomy — making recommendations, triggering actions, interacting with customers or

suppliers — what is the scope of that autonomy, and what are the escalation and override mechanisms? Integration is a period of elevated uncertainty; autonomous systems operating through it without recalibrated boundaries are a concentrated source of risk.

- Talent and capability assessment. Where does the AI expertise sit — in a central team, embedded in business units, outsourced to third parties? What is the retention risk, and what happens to model maintenance and governance if key people leave during the integration? AI talent is among the most mobile in any organisation; the integration plan must account for this specifically, not as a subset of general retention planning.

The intelligence layer is not a technology problem delegated to the IT workstream. It is a strategic integration risk that touches legal, commercial, operational, and regulatory dimensions simultaneously — and it moves faster than any of them.

5. The Chapter That Was Missing

We have spent three decades refining the M&A; integration playbook, and the result is genuinely good. The discipline around Day-1 planning, the rigour of synergy tracking, the hard-won lessons about culture and communication — all of this remains essential and none of it is diminished by what AI adds to the picture.

But the playbook was written for a world in which the assets being integrated were, between decisions, largely still. That world is receding. The intelligence layer is active, fast, and consequential — and it does not wait for the integration team to catch up.

The chapter that needs writing is not a footnote on technology. It is a fundamental extension of integration due diligence to cover assets that execute, learn, and decide — and that keep doing so whether or not anyone has told them the company just changed hands.