



# Human+Machine

Digital Magazine

---

**Issue #1**  
April 2026





---

## Contents

03 CEO Letter

---

05 Building an AI-ready data strategy:  
five questions CEOs should ask

---

13 How to lead the human side  
of an AI transformation

---

19 AI for Construction:  
what leaders should know now

---

23 Consumer marketing: how to drive AI  
value in a fast-changing consumer sector

---

29 AI in Private Equity:  
Deal-ready AI due diligence

---

37 Briefing: agents that work

---

# Welcome

## To the first edition of Human+Machine.

I came away from the World Economic Forum meeting in Davos earlier this year with a stark conclusion: organizations are struggling to realize the much-touted value of artificial intelligence (AI). AI is on everyone's agenda and pilots are manifold. But **scaling early initiatives** still feels hard.

That gap—between the technology's potential and today's state of play—is **why we have launched Human+Machine**. The aim is to regularly **share insights and practical advice** from our work in the field with leaders who are determined to capture AI's competitive advantages.

**This first edition of Human+Machine** examines two issues that are central to whether **AI investments pay off**—an organization's data strategy and its change management program.

**An AI model is only as good as the data upon which it is built.** And while most organizations have data aplenty, it's often inaccessible and unreliable, which leads to poor model performance. And when employees stop trusting models, adoption stalls. The article **"Building an AI-ready data strategy: five questions CEOs should ask"** argues that data is a mission-critical asset, making the development of a data strategy a leadership responsibility. Importantly, it also suggests how to craft that strategy in a manner that won't overwhelm the organization.

**"How to lead the human side of an AI transformation"** explains why people management is also critical to successful AI initiatives—a theme close to our hearts at Quantum Rise. Our firm belief is that organizations that make most progress understand that AI's value lies in the coming together of humans and machines. This article makes clear the extent to which AI will change people's work—the skills they need, what they will be held accountable for, and how roles will shift—and hence why change management programs are so important if new ways of working are to take hold.

The next pieces in the collection narrow the focus to look at three industries. **"AI for Construction: what leaders should know now"** demonstrates how AI can help owners and contractors get started with an AI playbook.

**“Consumer marketing: how to drive AI value in a fast-changing consumer sector”** notes that leaders might be tempted to go slow on AI initiatives when so much is changing and uncertain, but that would risk them falling behind competitors. Better to act strategically in a way that will deliver AI value today while laying firm foundations for bigger changes to come. The article takes stock of the way in which AI is upending consumer marketing. Workflows and team skills will need to change and a wholesale transformation of the operating and organizational model might ultimately be required. At the same time, the traditional customer journey is being disrupted.

**“AI in Private Equity: Deal-ready AI due diligence”** turns to the world of private equity (PE), arguing that AI is a key driver of value and risk for any company, but one rarely methodically assessed in PE deals.

The final article takes stock of new developments in AI and how intuitively easy is to use agents in 2026.

**Thank you for reading this first Human+Machine.** I hope it gives you a clear and useful view of where the **value of AI** lies hidden and how to begin to extract it.



**Alex Kelleher**

Founder and CEO, Quantum Rise



Alex is a multiple-time AI software and consulting entrepreneur. He was a founder at Vivid Edge (web consulting, sold to Framfab), Touch Clarity (machine learning optimization, sold to AdobeOmniure), and Cognitive Match (sold to Magnetic), and a leader at Magnetic (sold to Deloitte), Deloitte Digital (Global CMO) and Advantage Solutions (Chief Data Officer). He deeply believes in the transformative combination of human and machine, leading him to founding Quantum Rise together with business-services investment leader, Erie Street Growth Partners. Alex has a BA Hons in Experimental Psychology from Oxford University.



# Human+Machine

## Digital Magazine

---

**Issue #1**  
April 2026

---

**Louise Keely**

# Building an AI-ready data strategy: five questions CEOs should ask

*An organization's data strategy is key to whether AI investments pay off, making its development a leadership responsibility.*

Artificial intelligence (AI) looks set to become a fundamental part of business infrastructure, embedded in every work process. Competitive advantage, productivity, and growth increasingly depend on whether an organization can successfully, safely, and repeatedly adopt AI at scale. Yet multiple studies report that **70–95% of AI initiatives fail to make it to production**. We've seen many organizations begin to develop AI solutions that they expect to generate millions in value. But months or even years later, the solutions still haven't been deployed at scale, let alone yielded the hoped-for results.

The stumbling block is rarely the models themselves. Sometimes, it's **due to underinvestment in training and change management**, as AI deployment changes how decisions are made, how work is executed, and how risk is managed across the enterprise. But a fundamental problem lies in the weakness of the data upon which many models are built. **Data that is not well-structured** or labeled, incomplete, or insufficiently updated **leads to poor model performance** and business outcomes. And when employees stop trusting the model, adoption stalls and value evaporates. Robust data is therefore a foundation upon which trusted models depend and developing that data is a strategic capability that determines whether AI investments secure sustained competitive advantage.

The good news is that **CEOs do not need to be data architects or machine learning engineers** to craft a data strategy that delivers.

They do, however, need to recognize that a data strategy demands leadership attention because it's part of a business and operating model change, not a technology project. **Data is a mission-critical asset that needs to be managed as rigorously as any other**. That said, with constrained resources in mind, there is a pragmatic way of doing so that won't overwhelm the organization.



## ? The questions that matter

There are **five questions** CEOs should be asking themselves and their teams in pursuit of a data strategy that delivers:

1.

**Can we articulate the business outcomes and corresponding AI use cases the data will support?**

A data strategy determines how an organization sources, prepares, governs, and uses data. But this can only flow from the lynchpin of the strategy – knowing the data needed to drive specific business outcomes, be that using AI to automate parts of talent acquisition, providing **AI-guided assistance to customer service agents**, developing personalized marketing content, or offering pricing and negotiation guidance to procurement teams.

Many corporate leaders do not start at this point, jumping first instead to technology choices or governance. Both are important components of a data strategy. But they should be rooted in a plan designed to capture the business value of data. A CEO should be able to identify where value lies and home in the specific cases that will release it. Teams can then determine the necessary data, its format, and how it will be processed and governed. If this alignment is missing, organizations often end up investing in **broad data infrastructure** without a focused value creation path of a few key use cases, leading to low return on investment. Data readiness should be treated as a prioritization and sequencing problem based on business goals, not one of scaling.

AI readiness relies not only on the volume of data held. **Volume and diversity reduce model overfitting and bias**, but accessibility and quality also determine deployability. But most organizations have plenty of customer, marketing, and operational data emanating from many different sources: customer relationship management, point of sale, web analytics, enterprise resource planning, manufacturing systems, sensors, logistics, and third parties.

Many corporate leaders do not start at this point, jumping first instead to technology choices or governance. **Both** are important components of a data strategy. But they should be rooted in a plan designed to capture the business value of data. A CEO should be able to identify where value lies and home in the

2.

**Is our data AI-ready?**

specific cases that will release it. Teams can then determine the necessary data, its format, and how it will be processed and governed. If this alignment is missing, organizations often end up investing in broad data infrastructure without a focused value creation path of a few key use cases, leading to low return on investment. **Data readiness should be treated as a prioritization** and sequencing problem based on business goals, not one of scaling.

Of more concern, therefore, is whether the data is **clean, accessible, and reliable**, as AI products are only as strong as the data they retrieve and learn from. Sound financial planning is unlikely to emerge from analysis based on multiple different versions of a dataset, for example, while legal risks run high if contracts are drawn up on incomplete or inaccurate model data.

AI-ready data should therefore be of **high volume and high quality** in terms of its accuracy and completeness. But it should also be labelled and structured for AI model usage, transparent and with clear lineage, and automated.



### Labelling and structuring

Poor labelling and structuring undermine model performance, particularly in supervised learning, which depends on mapping inputs to clearly defined outcomes. Consider a retailer testing product promotions across standard shelves and end-of-aisle displays. Unless placement is properly labelled, the model cannot determine whether sales lift comes from the promotion itself or from premium positioning and therefore cannot optimize tactics. Structuring is equally critical for generative AI working with unstructured data such as contracts, policies, or transcripts. Effective retrieval requires consistent metadata schemas or embedding-based indexing to ensure models surface the right information.



### Transparent with clear lineage

Teams must understand the source of data and how it has since been transformed, for example if data is clustered in a specific way or summarized over different time frames. If the methods applied for these changes are documented, then they are easily understood and can be accounted for. Without this, anomalies can go unexplained. Likewise, clear, codified data lineage ensures its regulatory compliance, verifying, for example, that the use of customer data doesn't conflict with privacy laws.



### Automated

An automated data pipeline collects, cleans, standardizes, and distributes data, reducing human error and lag. Automation typically relies on scheduled ingestion processes, validation rules, transformation scripts, and monitoring tools that flag anomalies as they arise. This ensures that models and reports are fed consistent, up-to-date information, making AI outputs reliable and scalable and not dependent on manual data preparation.

Unfortunately, much of an organization's data today is unlikely to meet the required standard. This inevitably breeds a lack of trust in AI models. Teams feel obliged to double-check outputs, sometimes repeat the work or abandon the tool altogether. Yet the fix isn't necessarily to clean up all the organization's data and house it centrally in a data lakehouse, an expensive and lengthy process. A more targeted and pragmatic option is to prepare only the data required to feed an initial set of value-driving use cases. **A sequenced approach should be taken**, so that data investments build on one another, support evolving business AI priorities, and as data organizations can adapt as data requirements evolve with AI capabilities.

### 3.

#### Is our data governance a help or hindrance?

In many organizations, employees view governance as a bureaucratic layer of interference. Employees at one company we worked with viewed data governance as so restrictive (and poorly enforced) that they began building their own datasets and using their own software and engineers to manage and deploy the data.

Effective governance, in contrast, makes clear who is accountable for data quality, puts security, privacy, and compliance measures in place, and sets out the organization's ethical stance on the use of data and AI. Done well, it speeds execution and improves outcomes.

In our experience, effective governance is best achieved by setting up a **two-body model**: a working group that designs, and a council that enforces and arbitrates.

The working group defines how data should be governed and **is responsible for compiling documents such as a governance charter**, a domain inventory, ownership matrices, policy drafts, quality metrics, audits, and readiness reports. It's important that the work is conducted by those close to the data and business needs, so the group might be headed by a business unit director or senior manager but include data managers from various functions such as finance and operations, a data architect who is responsible for the structure, content, and labeling of the data based on the business needs, and compliance liaisons who can advise on their areas of expertise, such as legal, risk, and human resources.

The data governance council approves decisions made by the working group, resolves conflicts, and enforces standards across the enterprise with a view to increasing trust in AI and accelerating execution. Governance responsibility must sit alongside enforcement authority, otherwise the council will be ineffective. The data governance council is likely to consist of senior leaders

from key functions such as finance, IT, operations, risk, and change management.

These are all part-time roles; members of both groups maintain their current positions whilst serving in these groups. Governance should reduce friction and increase organizational trust, not add control layers that slow down progress.

## 4.

### Do we have the capabilities needed for data readiness?

Many CEOs underestimate the depth and mix of capabilities required for sustained AI deployment, including around data. **There are three challenges that many companies have:**



Talent scarcity and cost: the need for people who can build data foundations are widespread and the skills are highly specialized.



Rapidly evolving skill requirements: the requirements for AI data readiness can be different from other work, and changing AI capabilities itself requires regular upskilling.



Misalignment between business and technical teams: when liaisons to translate between the two groups and to voice and resolve concerns or conflicts is insufficient.

Therefore, key roles include:



Stewardship and quality management. These roles help ensure data quality.



Technical support. This includes data engineers, data scientists, and AI and machine learning engineers who manage and govern the data effectively. Each of these groups have different roles. Data engineers build data pipelines, data scientists analyze data and develop models to answer business problems, while ML engineers productionize and scale models.



Business translators to define requirements and ensure adoption. These collaborate with technology and data teams and business users to ensure that the right data and tools are prioritized and the user experience good.

It's also important that employees across the organization are comfortable using data within AI tools, which will likely require training, reskilling, and change management programs. **Capabilities must be embedded in business units,** not centralized exclusively in a technology organization.

## 5.

**Are we doing enough to establish a data-centric culture?**

Even strong data strategies, those grounded in clear business goals, AI-ready data, effective governance, and the right capabilities, will fail if tools are not adopted in **day-to-day work**. Scaling AI requires a **shift** to a genuinely data-centric culture, where decisions are expected to be evidence-based and where shared, trusted data replaces locally managed spreadsheets and informal analyses. Employees must be comfortable acting on model outputs and documented insights. Leaders must reinforce that expectation consistently and practice it themselves. This demands more than new tools; it requires changes to incentives, workflows, and accountability structures so that data-driven decision making becomes standard practice rather than optional behavior.

Leadership plays a central role in embedding these norms. Incentives can be tied to AI adoption, measurable productivity improvements, and explicit ownership of data quality within business domains. Core decision processes, such as pricing, hiring, procurement, or risk reviews, can be redesigned to require reference to shared dashboards or AI outputs, with clear documentation when recommendations are overridden. Named data owners, transparent quality scorecards, and defined escalation paths reinforce accountability. At the same time, **structured AI literacy programs** and visible executive use of AI tools signal that data-driven execution is expected at every level. Change management, training, and workflow redesign are therefore not ancillary activities; they are integral to realizing value from AI investments.

## The questions that matter for AI data strategy



AI transformations succeed not because organizations experiment with models but because they build data foundations that allow those models to operate reliably at scale. Alignment on value creation priorities, AI-ready data, effective data and AI governance, the right capabilities, and cultural norms that support data-driven decisions are mutually reinforcing elements of an AI-ready data strategy and an operating model shift. When one is missing, the shift will stall. When all are in place, AI can be deployed across the enterprise rather than episodically.

Organizations that attempt to “fix all their data” before pursuing AI will often stall under the weight of large, unfocused transformation programs. Those that take a sequenced, use-case-driven approach build momentum and demonstrate value more quickly. The cost of delaying this discipline compounds over time: as competitors strengthen their data foundations around priority use cases, the execution capability gap widens. **To get started, if your organization has not already done so, identify three high-value AI use cases and assess**, within the next 30 days, whether your current data quality, accessibility, governance, and ownership structures are sufficient to support them. The goal is not perfection across the enterprise, but clarity on where readiness exists and where targeted investment is required.



### Louise Keely

Managing Director, Consumer



Louise is the Managing Director for the Consumer Vertical at Quantum Rise. She brings deep expertise as a growth strategist and data scientist, and as an advisor to senior executives across consumer sectors and the data and analytics firms that support them. Previously, Louise was a partner at Bain and EY-Parthenon, and served as EVP of the global retail practice at Nielsen. Louise holds a PhD in Economics and an MSc in Econometrics from the London School of Economics, where she was a Marshall Scholar, and a BSc in International Economics from Georgetown University.

---

**Yingyang Wu**

# How to lead the human side of an AI transformation

*AI transforms the way people work, which is why scaling its adoption depends so heavily on a carefully constructed change management program*



Investment in AI has grown quickly but returns are uneven. Research last year by the Massachusetts Institute of Technology found that 95 percent of AI initiatives worth some \$30 to \$40 billion failed to generate measurable returns, never reaching scale<sup>1</sup>. The remaining 5 percent made handsome gains, however.

**AI tools impact people  
and the way they work  
compared with older  
technologies**

The same research points to the cause of such a wide divide. The most frequently encountered barriers to scaling pilots were an unwillingness to adopt new tools and organizational blockers such as inadequate executive support or change management programs.

In short, scaling AI is primarily a people challenge.

That's hardly surprising given the extent to which the value of AI is ultimately capped by the scale of its adoption, such as ERP and CRM, and it explains why getting change management right is so important. AI tools are only valuable if adopted, and they will only be adopted if people understand the significant changes afoot, if senior leaders back the changes, and employees are fully supported to make them.

<sup>1</sup> The GenAI Divide: State of AI in Business 2025, July 2025

## AI's impact on people

**AI impacts people and their work in three key ways. New skills will be required, there will be new accountabilities, and job designs will change. All can hamper the adoption of AI if not addressed.**

**Skills.** The introduction of new technology often demands a degree of change. People may need to perform their roles a little differently with slightly different processes, and they may need to become proficient with new systems. But the changes are usually minor compared to what AI demands. AI demands that employees deploy new and often higher-level skills. For example, a front-line employee who once followed a rigid process must now think critically about how to prompt an AI tool, evaluate work outputs, and take responsibility for the eventual results, all of which are usually associated with managerial roles. So, although the mantra from AI enthusiasts is that AI will handle the mundane work so that people can focus on more interesting problems, that transition will only happen with a significant upskilling of the workforce.

**Accountability.** Generally, the introduction of a new system doesn't change an employee's accountability much. It may alter the click path for a common task, but the employee remains accountable for certain basics over which they have control, such as inputting the right data or making the right calculations. With AI, employees essentially become accountable for decisions made by a machine that they didn't build and whose internal workings aren't clear. This creates new pressures. People worry about being blamed for an AI mistake or feel forced to choose between moving fast and double-

checking everything. There is new work to be done too. For instance, although AI outputs look polished, employees can't just skim for typos. They must validate the outputs, which may mean establishing an audit trail that documents what was asked, what the tool suggested, and why a decision was made. And although AI generates output in seconds, it doesn't always consider the potential risks of its conclusions. Employees must make that assessment. Hence, if an organization doesn't establish guardrails—what good verification looks like, when escalation is required, and who signs off on the final output—using AI can feel like extra work with extra risk, discouraging its adoption. Alternatively, in the absence of guardrails, employees might lean too heavily on AI without sufficient verification. Both scenarios increase operational risk and can slow a transformation.

**Job Design.** New skills and new accountabilities reshape roles. For example, a customer support agent might spend less time composing responses to customers and logging cases and more time diagnosing intent, choosing the right resolution path, validating policies and the account context, and deciding when to escalate matters. It becomes a very different job, requiring broader judgment and more cross-functional handoffs than a scripted workflow, and one that carries more responsibility. If organizations don't update job designs accordingly,

employees can experience the change as confusion or mismatch of what their job is. Moreover, accurate job designs lay the foundation for a

new talent strategy to ensure hiring, training, promoting, and performance management practices are aligned to the new expectations.

## A change plan that manages AI's people impact

**All this explains why change management is so important for realizing AI's value. A change management program that encourages adoption will likely have three main components: a clear and compelling vision of the changes required, a strong change coalition at the top, and strong support for employees.**

### **A clear and compelling vision.**

Traditional change management starts with a strong business case and a clear future-state vision. With AI, you'll move faster if you get specific about what is changing: the skills, accountabilities, and roles. If you don't, teams stall in pilots, resist new workflows, or route around the tooling. So, stand up an operating model workstream to redesign workflows and roles.

### **A strong change coalition at the top.**

AI deployment can affect employees at all levels of an organization. That's why it's important to gather the support of senior leaders and heads of functions as much as that of more junior staff. Those leaders might be wondering whether their function will be seen as less valuable in a world where AI can perform much of the work, whether they will have less autonomy with decision rights shifting to those designing the AI tools, or whether their teams will inherit new liabilities as AI enters the workflow. Such concerns need to be addressed directly. Ignoring or suppressing them will likely brew disquiet, with leaders denying approvals, withholding resources, limiting the scope of initiatives, and raising exceptions that block standardization. AI adoption within an entire business could be at risk.

### **Support employee adoption.**

The change management framework for encouraging the adoption of AI within the workforce is the same as for any other transformation, namely building a shared understanding of what's afoot, building trust, building ability, and reinforcing the new ways of working. The detail of the program addresses the specific challenges of AI initiatives, however.

- **Build a shared understanding of the layers of change.**

Change management programs typically emphasize the business case for change, the risks of not changing, and the future-state vision. Those elements still matter for AI initiatives, but the changes are so significant that employees require more detail. Without it, they are likely to construct their own narratives in response to news that the organization is implementing AI—a narrative often shaped by anxieties about job security, the potential for AI mistakes, and shifting expectations. To avoid this, ensure that leaders and managers translate the organizational vision that defines what is changing and why to one that resonates with people lower down in the organization.

- **Build trust by addressing concerns head-on.** As with organizational leaders and heads of functions, employees' concerns about AI need to be addressed if they are to trust the technology. A well-tested formula for building that trust is empathy + an action plan + consistent follow-through. So, if employees are worried about their jobs, acknowledge their concerns, provide reassurance and a workforce redeployment plan that explains where the saved time will go. If employees fear AI mistakes, explain the guardrails, such as where the use of AI is permitted, what needs verifying, and how to escalate uncertainty. And if trust in outputs is low, demonstrate how recommendations are produced and where human oversight is required. In all such situations, provide visible follow-through so employees can see that raising concerns leads to concrete changes.
- **Build knowledge and ability.** Traditional change management focuses on procedural and system training, whereas employees often need to learn higher-level cognitive skills when adopting AI. So, start by identifying the new skills required, then help employees acquire them through formal training, working sessions, coaching, reference materials and more. For example, when an AI-enabled dashboard is introduced, a business analyst's job shifts from building spreadsheets and formulas to translating data into insights, making recommendations, and influencing decisions. Support here should include helping the analyst not only learn how to navigate the dashboard but also how to interpret signals, test assumptions, surface risks, and turn insights into clear actions. Shared standards such as common approaches to prompt patterns, validation steps, and documentation norms can help, while communities of practice can accelerate learning if teams share their experiences, good and bad.
- **Reinforce new ways of working.** To sustain change, organizations will need to establish systems that help make the new behaviors stick. Feedback is key. Leaders and managers should share success stories, coach verification habits, and recognize teams that adopt good practice. They should also review how adoption is progressing at the company or department level, looking at AI usage in target workflows, verification and escalation rates, quality outcomes, and cycle-time improvements. Such metrics can show where AI-enabled workflows might be stalling and help fix the cause. Leaders should also provide plenty of opportunities for employees to check in and refresh their knowledge, including drop-in sessions, short refresher courses, and communities of practice.

The bottleneck to adopting AI at scale is rarely access to models. Rather, it's the organization's ability to change how work gets done. To manage that change, leaders will need to understand the new skills required, the new accountabilities, and how roles will evolve. Without this framing, the result is likely to be an accumulation of pilots and, at best, pockets of productivity wins. With it, leaders

can build an effective change management program that encourages the adoption of AI by making clear what lies ahead, building a strong coalition at the top, and supporting employees for what the future holds. That is what moves an organization from experimentation to creating value at scale, and it is how companies land on the right side of the AI divide.



### Yingyang Wu

Director of Client Training and Enablement



Yingyang Wu is the Director of Client Training and Enablement at Quantum Rise. With over a decade of experience in corporate learning and development and organizational change, Yingyang has led large-scale initiatives including enterprise technology transformations, organizational restructures, and workforce development programs across industries such as healthcare, supply chain, marketing, and professional services.

---

**Hugh Seaton**

# AI for Construction: what leaders should know now

*Artificial Intelligence is the largest investment project in history; US data centers alone will top \$1 trillion in cumulative spending in 2026. That investment has produced real capability, and the pressure to adopt AI has reached every construction boardroom.*

## AI is Not Another Application

Contractors who spent 20 years learning to assess and integrate software applications are finding that AI requires a fundamentally different approach.

The playbook for adopting a new application does not work for a technology that can produce content, make decisions, and act within your processes.

Many AI projects treat the technology as if it were another point solution: contractors identify a problem, commission an AI tool to solve it, and deploy. This approach may yield a working tool, but it builds no internal capability to evaluate, manage, or expand AI use over time.

AI is a general technology, like the cloud or mobile, and like those earlier waves it will reshape how work gets done across the organization. What makes AI distinct is that it can produce and act. **An AI agent can draft an RFI response, check it against the spec, and route it for review.** That makes AI an actor within a process, not a passive tool used by one. **No other technology has had this capability**, which raises a question every contractor will need to answer: where does AI fit in our organization and its processes?

## AI is accelerating

The pace of AI improvement has gone from remarkable in 2023 to astonishing in 2026, and it is still accelerating. What was impossible months ago is commonplace now, and this pattern will continue. For context, it took the web 20 years to evolve from browsers to enterprise SaaS; we are roughly 4 years into that cycle. Fundamental innovations still ahead will expand the reliability, scope, and scale of what AI can do. Buying a

specific AI application today means investing in a moving target. A better use of that investment: **build your people's ability** to evaluate AI, and map your processes clearly enough that you can rework them through ongoing automation as the technology matures.

Answering that question well requires understanding **where AI capability actually stands today, and how fast it is moving.**

Today's **AI splits into two modes** that matter for construction. A **copilot** assists an individual; it helps a project engineer draft a letter or summarize a long spec section. An **agent** executes steps within a workflow; it can pull data from your ERP, compile a daily report, and flag exceptions for human review. The copilot needs a skilled user.

The agent needs a very well-defined process.

Many contractors have started with copilots, which is

sensible, but the larger gains will come from agents, and agents demand process clarity that most organizations do not yet have.

Before approving an AI agent pilot, confirm the target process meets these criteria:

### 1. Highly repeated.

Daily reports, submittal logs, and pay application line items all recur on every project.

### 2. Consistent set of steps.

Submittal review follows a predictable sequence: check against the spec, verify product data, route for approval.

### 3. Moderate complexity.

Formatting, completeness checks, and references to standards. RFI preparation, for instance, requires pulling from drawings and specs but follows a known template.

### 4. Consistent set of sources and integrations.

The data lives in known systems (your ERP, project management platform, or document control tool) rather than scattered across emails and texts.

These characteristics apply to many of the support and preparatory tasks that, while essential, consume time that could go toward putting work in place. The challenge is fitting AI into these processes so it saves time and expands what is possible, without creating new exposure. An agent that drafts an RFI response, for example, could introduce a scope interpretation the PM never intended; an agent that generates daily reports could miss a safety incident buried in a subcontractor's notes. Managing these risks requires understanding the process well enough to know where human review must stay in the loop.

**Selecting the right processes for AI integration requires understanding them as mappable workflows, not tribal knowledge.**

Construction, **like most industries**, does not normally document most processes clearly. Workers have a task and a rough set of steps, and they often figure it out fresh each time. This is one reason highly repeated tasks are the right starting point for agents; the steps tend to be consistent and well understood by the people doing the work.

### Start with People and Process

AI will continue to grow in capability and breadth of adoption, creating pressure to apply it to more and more of the business, much as mobile and cloud software did before it. Preparing for that pressure means building the organizational **ability to adopt, integrate, and succeed with AI** over time.

**That starts with understanding your processes well enough to prioritize; which ones are genuine pain points, and which of those pain points match what AI can actually do today.**

The people best able to map a process reliably are those currently executing it. Training in AI use, beyond basic prompting, makes this kind of ground-level analysis possible. When a project engineer understands what an agent can and cannot do, she can identify where automation fits in her submittal workflow and where it would introduce risk. That knowledge contributes to both adoption and better design of AI solutions as they are built.

Early results bear this out. Many contractors report that their 2025 AI pilots underdelivered. Some of that was technology immaturity, but the more common problem was treating a pilot like adopting a new application, with too little training and no real mapping of AI capability to process type. **Combining early training with process analysis**, so the right process is selected and the solution fits how work is actually done, is the **most reliable way** to avoid repeating those outcomes.

## **Winning Early**

None of this upfront work needs to be extensive, and it should not slow you down. Pick a bounded internal function, say accounts payable or quality assurance, and spend two to three weeks analyzing its processes, steps, and data sources to **identify the best candidate for an AI agent**. That focused scope limits risk while building the internal muscle to evaluate AI for broader use. A contractor who runs this exercise on pay application reconciliation, for example, will learn more about where AI fits in their organization than a year of vendor demos could teach.

**AI's scope and pace are unlike any prior technology wave.** Contractors who invest in process clarity and people development now will be positioned to absorb each new round of AI capability as it arrives, compounding the advantage over time. **The first step is modest: choose one process, map it, train the team, and build from there.**



**Hugh Seaton**

Managing Director, Construction



Hugh is the Managing Director for the Construction Vertical at Quantum Rise. He started his construction technology innovation journey translating specifications into Chinese in Taiwan in the 1990s, and has been GM at the Construction Specifications Institute, founder of the AEC Hackathon, CEO of a construction software company, and now is leading the construction vertical at Quantum Rise, to help this essential industry transform with AI.

---

**Louise Keely**

# Consumer marketing: how to drive AI value in a fast-changing consumer sector

*AI is upending today's marketing function, though the end game is unclear. Don't let uncertainty about the future prevent you from realizing AI value today.*

The marketing department has been the first to experiment with AI in many US consumer companies. With the average B2C company spending close to 10 percent of revenue on marketing<sup>1</sup> and middle market companies often spending considerably more, AI's potential to cut costs while boosting results has been greeted enthusiastically<sup>2</sup>. Many companies today use AI to generate content, analyze social media data, and test campaigns, for example.

A few among them have already earned substantial value from their efforts (see sidebar, "AI in action"). Many others, however, are stuck in the experimental phase without a strategy for progressing. CMOs and CEOs frequently tell us their teams are testing various AI tools and can list what they are. But when asked about governance processes aimed at using data efficiently and minimizing risk, training programs, or plans to revise workflows and team structures in response to AI's capabilities, they often struggle to give answers.

**AI tools are developing rapidly, making the end game uncertain. Workflows and team skills will need to change.**

Additionally, they frequently will have no idea of the value created from using those tools.

In their defense, marketers have a lot to contend with. AI tools are developing rapidly, making the end game uncertain. Internally, workflows and team skills will need to change and a wholesale transformation of the operating and organizational model might ultimately be required. At the same time, the traditional customer

journey is being disrupted. As shoppers increasingly use AI tools and agents to search, compare prices, and purchase on their behalf, how

will marketers market? Certainly not as they do today.

One response might be to do little until the future is more certain, especially if budgets are tight. But that would risk an organization falling behind competitors. Better to act strategically in a way that will deliver value today while laying firm foundations for bigger changes to come.



<sup>1</sup> The CMO Survey, Duke Fuqua School of Business, 2025; CMO Spend Survey, Gartner, 2025.

<sup>2</sup> The figure is higher in some sectors. CPG and media and communications companies, for example, spend around 25 percent of their revenue, on average.

## AI tools are becoming increasingly powerful, calling into question marketers' role.

Some things are already clear. First, AI will impact each stage of marketing, from market analysis to campaign activation and performance tracking. Second, AI tools will become increasingly powerful, augmenting and eventually automating many of the workflows that marketers currently execute. **See Appendix for examples.**


Today, AI tools mainly assist humans in their work. They can help develop campaigns, but humans remain in the loop, approving every step in the process including

creative output, audience and targeting decisions, and budget allocation, to name a few. The assistance can come both in the forms of assisting in manual tasks such as allocating budgets across marketing channels as well as creating more tailored, even personalized, content, that should improve consumer response.

**AI agents should eventually be able to take full control, coordinating and executing multiple steps in the campaign process.**

Over time, however, AI is expected to execute campaigns autonomously. Humans might remain “on the loop” for years, meaning they still monitor, guide, and intervene as needed. But AI agents should eventually be able to take full control, coordinating and executing multiple steps in the campaign process, with input from other AI agents. An AI agent will be able to rapidly launch a campaign with numerous channel variations, personalize and adapt it dynamically, and incorporate insights from previous

campaigns to improve design and execution. Only a fraction of marketing organizations are already using AI agentically to automate certain workflows (not yet end to end), and they tend to be larger ones.



All this raises questions about marketing's future. For example, if AI can create and implement models for allocating marketing budgets, continually refine the strategy, run experiments to test campaigns, analyze the results, and make recommendations, what is the role of the quantitative marketing analyst? Someone still needs to build, maintain and oversee AI agents. Will that responsibility sit within a marketing function or be outsourced to AI experts with quite different skills? Will certain types of marketing work remain with humans, even if AI could technically do it, because there are aspects of human creation or judgement that AI is unable to fully replicate? Or will a different organization of customer, related decisions emerge, entirely distinct from what exists today?

**Marketers also need to confront how AI is changing shopping behavior.**

Marketers also need to confront how AI is changing shopping behavior. In a

world where AI-powered agents can research, choose, and buy a product for consumers, even anticipating needs before they arise, the traditional customer journey and marketing funnel will likely become obsolete, at least for certain types of purchases. Marketers need to address consumer intent and market to the decision maker. But whom should

they address in a world where, for instance, an AI solution orders groceries based on a picture of the contents of a consumer's

pantry and a record of their previous purchases and preferences? For at least some types of purchases and shoppers, and stages of the shopping journey, marketers find themselves marketing to large language models and AI agents, at least as an intermediary if not the outright decision maker. This change is already underway.

## The way ahead

**So much change has left CMOs and CEOs** we speak with unsure as to whether and how they should press ahead with AI initiatives within their current marketing organization and operating models or embark on a transformation of those models, reimagining the target audience, the solutions and processes required to reach and influence the audience, and the organization, skills and workflows of a marketing team. **Our view is that an extensive reorganization, or even a less comprehensive one than described above, is not the right starting point,** not least because it's unclear what the optimal model might be, as AI tools and capabilities are still evolving. **A CMO's priority should therefore be to put in place the fundamentals that support value creation today and tomorrow.** There are four no-regret moves that will help CMOs and their teams make decisions on how and where to deploy AI in their organization in support of value, while simultaneously laying a solid foundation for marketing's evolution.

1.

**Learn and educate.** There is still much hesitancy over AI deployment. Some firms still don't allow its use, concerned about the risks, while many employees don't know how to use AI in their work or the extent of its utility. Others avoid it, concerned about its impact on jobs. Training programs are therefore needed to help teams gain confidence and adopt AI in their work. At a minimum, marketing leaders should introduce AI-enabled productivity tools that help with CRM, email marketing, content creation, and more.

3.

**Assess AI readiness.** Ambition must match capabilities. So decide whether you have the data, technology and people you'll need, testing prioritized use cases for constraints and taking action to remove them. Check data quality, access, governance, and reuse across key domains. Review the technology stack for scalability, security, and AI delivery. Assess skills, roles, and operating model support.

2.

**Set a vision and priorities.** Set an ambitious goal for what you might like AI to achieve - a 30 percent reduction in marketing costs within five years, for example. You can't devise an overly detailed plan for achieving that goal. Too much will change in the intervening years. But you can draw up an initial roadmap with short and longer-term initiatives, prioritized around business value and feasibility.

4.

**Build pilots and get them in the hands of users, early and often.** Start with a workshop to define target outcomes for a use case, map current and future workflows to identify and address bottlenecks, and design the path to build and deploy (including any required integration with legacy systems). Next, validate rapidly with a prototype built with product-ready code, and include from the start lightweight governance, such as adherence to brand guidelines. Then test with real users before building and deploying the product at scale. It's a derisked process that includes the user in the development process to validate feasibility, user need, and value at stake in weeks, not months.

\*\*\*

Important organizational and operational changes will of course have to be made in the next few years to accommodate the changes that AI is unleashing within the marketing function of B2C companies of all sizes. But the starting point is not a major overhaul. Rather, begin by deploying AI within current work routines and processes, guided by a roadmap and metrics that ensure efforts drive value. The preparation described here will stand any organization in good stead for bigger changes to come. You need strong foundations and experience with AI on which to build a marketing future, whatever that might hold.

## AI in action

Some marketing organizations are already deploying AI to create value. Here are a few examples of those with which we've worked.

### Governance

A brand agency needed a faster, more reliable way to apply brand standards. Wading through lengthy brand books to check the rules and reviewing their application not only slowed creative production but raised compliance risk. We helped build an AI-driven brand governance platform that ingests company brand guidelines (PDFs and related assets) and makes them searchable. Now, through a chat-based interface, users can quickly discover what's required and validate work. The result: brand-compliant content delivery time cut by more than 50 percent, costs by two-thirds, and a doubling of project margins, all while improving brand governance consistency across customers' campaigns.

### Measuring performance

A group of regional U.S. retailers were struggling to demonstrate the impact of advertisers' campaigns on their media networks and failing to capture advertising spend from larger players as a result. Quantum Rise worked with a global data science company to

develop a standardized retail media health check: a platform that scores performance against various metrics, including against category leaders, and suggests data-driven improvements. The outcome: retailers can consistently demonstrate ROI potential, helping to attract more advertising.

### Audience building

An advertising services company needed to improve its data capabilities after an acquisition to improve its location-based audience-building and targeting capabilities. Inconsistent regional taxonomies, lack of international coverage, and various data storage and labeling inconsistencies were making the task hard. With the use of AI, Quantum Rise built consistent regional boundary definitions across databases and implemented standardized data labeling, and deployed mapping and visualization AI-enabled tools to better display targeting recommendations. The result was more accurate targeting of consumers by location that took advantage of its expanded geographic coverage.




### Louise Keely

Managing Director, Consumer



Louise is the Managing Director for the Consumer Vertical at Quantum Rise. She brings deep expertise as a growth strategist and data scientist, and as an advisor to senior executives across consumer sectors and the data and analytics firms that support them. Previously, Louise was a partner at Bain and EY-Parthenon, and served as EVP of the global retail practice at Nielsen. Louise holds a PhD in Economics and an MSc in Econometrics from the London School of Economics, where she was a Marshall Scholar, and a BSc in International Economics from Georgetown University.



---

**Fabio Ferraretto, Josh Kohn**

# AI in Private Equity: Deal-ready AI due diligence

*How private equity deal teams evaluate  
AI value drivers and risk vectors*



AI is rapidly becoming a **standard part of the acquisition process**, much like cybersecurity did more than 15 years ago. For private equity deal teams on both the buy side and the sell side, it is now a meaningful driver of both upside and downside in most sectors they touch. **AI can accelerate growth, improve margins, and sharpen differentiation.** It can also disrupt core offerings, erode moats, or expose hidden operational and governance gaps.

Yet most technology due diligence still focuses on systems, infrastructure, and security. It rarely looks at AI in a structured way or assesses how AI might change the company's economics and competitive position over the next three to five years. The result is a new blind spot in modern M&A.

A practical way to close that blind spot is to evaluate AI through two lenses that mirror how deal teams already think: value and risk. Rather than asking, "What AI does this company have?" the more relevant questions are, "Where does AI actually create or protect value?" and "Where does it introduce exposure that matters to our thesis?"

## AI Value Drivers

### Revenue

New AI-enabled products, smarter targeting and pricing, stronger retention.

### Profitability

Automation of manual work, fewer errors, better use of assets and inventory.

### Speed & edge

Faster decisions and iteration, AI embedded in workflows, harder-to-copy capabilities.

Through the first lens, **AI value drivers** show where AI can move the top line, improve profitability, and increase speed and differentiation. That might include AI-enabled products and services, better targeting and pricing, more efficient operations, and faster iteration that keeps the company ahead of competitors.

## AI Risk Vectors

### Where deals get exposed

#### Tech & data

Brittle architectures, weak MLOps, unclear data rights and provenance.

#### Execution & talent

Key-person dependency, manual processes, Shadow AI outside formal controls.

#### Strategy & governance

Platform and agentic AI shifts, regulatory pressure, reputational exposure.

Through the second lens, **AI risk vectors** reveal where AI exposes technical and data fragility, execution and talent risk, or strategic and governance concerns. Examples include brittle architectures, unclear data rights, key-person dependency, or the risk that platforms and agentic AI interfaces will reshape markets and distribution in ways that undermine the current business model.

## I The pillars of AI maturity

To move beyond impressions and war stories, deal teams need to break AI into components that can be tested with evidence. **One useful structure is to look at AI value drivers and risk vectors across five pillars of maturity:** product and intellectual property; data and machine learning operations; tech stack and DevOps; talent and ways of working; and governance and risk.

- In the **product and IP pillar**, the focus is on how central AI is to the offering and how durable that advantage might be. Does AI actually show up in the product or service in a way that customers value, or is it mainly marketing language? What, if anything, is truly proprietary — models, data, workflows, integrations — and how hard would it be for competitors or platforms to replicate?
- The **data and ML operations pillar** looks at the quality and readiness of the fuel that AI depends on. Here the questions turn to whether core datasets are complete, clean, governed, and clearly owned; whether they are structured to support current and future use cases, including real-time needs; and how reproducible and automated the model development and deployment processes really are. Strong data and ML operations make AI reliable and scalable. Weaknesses here often show up later as instability, high cost to serve, or an inability to move beyond proofs of concept.

- 
- **Tech stack and DevOps** determine how fast the company can evolve its AI capabilities. The underlying architecture needs to support the performance, security, and latency requirements of AI workloads, and CI/CD practices should extend to models and data pipelines, not just application code. The question is less about any specific tool and more about flexibility: can the company adopt new AI approaches or providers without major rework, or is it boxed in by past decisions?
- 
- **Talent and ways of working** bring the analysis back to people. A balanced mix of data science, engineering, product, and domain expertise is a starting point, but the real test is how AI work gets done. Are initiatives prioritized and tied to P&L owners and concrete outcomes, or are they largely experimental? Is AI woven into cross-functional teams, or confined to a small group of specialists? How much key-person risk exists around one or two individuals who hold critical knowledge in their heads?
- 
- Finally, the **governance and risk pillar** looks at how safely AI can be deployed and scaled. That includes the presence (or absence) of controls around privacy, model risk, security, and responsible AI, as well as how AI-related incidents are detected, escalated, and addressed. It also asks how prepared the company is for evolving regulation and changing expectations from customers, partners, and regulators. The upside here is the ability to scale AI with confidence; the downside is regulatory exposure, customer pushback, or reputational damage.

**Seen through these five pillars**, AI due diligence becomes a structured exercise. For each pillar, deal teams can identify the most important **value drivers and risk vectors**, gather **evidence**, and assess how they affect the case.

## What AI due diligence actually finds


**When this approach is applied in live deals, the findings are often both technical and strategic. A few anonymized examples illustrate the range of issues that can emerge.**

**In one case**, a retrieval-augmented generation chatbot looked impressive in demos and was positioned as a near-ready product. Underneath, it was essentially a monolithic script: no tenant isolation, limited logging, and synchronous ingestion that could block user traffic. The idea addressed a real need, but the implementation carried technical and data risk that would require meaningful re-engineering. The investment thesis had to treat it as a proof of concept, not a finished asset.

**In another engagement**, a computer vision platform showed strong model performance and recurring revenue. Deeper review found that training still relied on older hardware, many machine learning workflows were manual, and advanced R&D was concentrated in a small group. The value drivers were real; so were the risk vectors: technical debt, higher operating costs, and key-person exposure. Any case built on continued growth needed to account for the cost and timing of modernizing the environment and broadening the talent base.

**A third example** involved a subscription management product tightly integrated with a major commerce ecosystem. The product had a rich feature set and loyal customers, but emerging agentic AI interfaces and new platform APIs were making it easier for merchants and consumers to handle subscription events through native agents that bypassed the product's experience. The company's most defensible value drivers lay in capabilities that would be harder for generic agents to reproduce, such as deeper analytics and targeted retention channels. The investment plan needed to lean into those strengths and assume some erosion of advantages that the platform could standardize.

Across these situations, AI due diligence did more than confirm or refute technical claims. It surfaced architectural gaps, operational drag, talent concentration, and market shifts that materially influenced valuation, deal structure, and post-close priorities.



## Putting the lens to work: buy side and sell side

For **buy-side deal teams**, AI due diligence helps clarify how much of the thesis actually depends on AI-driven value drivers and how credible those drivers are. It highlights where risk vectors could affect base, upside, and downside cases, and what must be addressed in the first 90 days and over the next 18 to 24 months to protect and unlock value. It also informs questions of price and structure: which issues justify a change in valuation, and which can be handled through covenants, earn-outs, or targeted post-close investment.

On the **sell side**, the same structure can be used much earlier in the ownership cycle. A portfolio company preparing for a process can use the value driver and risk vector lens to identify AI strengths and articulate them clearly in the equity story and CIM, to find and address obvious AI risk vectors before buyers do, or at least frame them proactively, and to assemble a credible AI roadmap that aligns with likely buyer narratives and questions. In markets where AI maturity is increasingly a filter, being able to tell that story clearly is becoming a differentiator in itself.

The discipline is the same in both directions. What changes is when it is applied and how the findings are used.

### Five Questions Every Deal Team Should Ask

#### Five questions for the deal team

- 1** Where does AI really show up today, and how much does it matter to revenue and margin?
- 2** What proprietary data and workflows power AI, and do we clearly own and control them?
- 3** How reliable and scalable are the current AI systems, beyond what we see in demos?
- 4** Which AI-related risks could change our base, upside, or downside cases?
- 5** If we own this asset, what AI work is non-negotiable in the first 90 days and the next 18–24 months?

## From findings to action

For AI due diligence to matter, it has to show up where decisions are made and capital is deployed.

On the investment committee side, **AI findings can sharpen three familiar components** of the memo. **First**, they make the upside thesis more concrete by

spelling out where AI is expected to drive revenue growth, margin expansion, or multiple uplift over the hold period. **Second**, they clarify the dependencies and investment required by highlighting the most important gaps in data, tech stack, talent, and governance and estimating what it will take to close them. **Third**, they define the key AI-related risks that matter most to the case and how they influence base, upside, and downside scenarios.

On the post-close side, AI diligence can be translated into a concise execution view that fits a Consulting 2.0 mindset, one that emphasizes building durable, productized

capabilities rather than running disconnected projects. In practice, that often looks like a sequenced 90-day plan focused on discovery, quick wins, and non-negotiable remediation; a high-level 18 to 24 month AI roadmap aligned with broader value-creation work; and a short list of AI initiatives that are designed from the start to become ongoing products, not one-off experiments.

This **bridge from findings to action is where AI due diligence proves its value**. It shapes price and terms. It improves negotiation around representations, warranties, and covenants. It gives management teams a head start on the work that will matter most to both value creation and risk management.

## From Diligence to Action in Three Outputs

### What “good” AI diligence produces

#### Sharper IC story

A clear AI upside thesis plus the dependencies and investment it requires.

#### Concise risk view

The few AI risk vectors that actually move valuation, structure, or covenants.

#### Execution start line

A 90-day plan and 18–24 month AI roadmap tied to value-creation priorities.

## What “deal-ready” looks like

Deal teams do not need deep technical expertise in every AI technique, and they do not need to resolve every uncertainty about the future. What they do need are consistent answers to a small number of questions: **Where are the real AI value drivers** in this business, today and over the next few years? **What evidence supports or challenges the claims** being made about AI? **Which AI risk vectors are we taking on**, and are they acceptable given the price, structure, and plan?

A thoughtful AI due diligence process, built around value drivers, risk vectors, and a clear view of the five pillars of maturity, can provide those answers. It does not remove risk; **it makes that risk visible and manageable**. In a market where AI is increasingly embedded in every deal, that is what it means for private equity teams to be truly deal-ready on AI.

## Five pillars of AI maturity

### Product & IP

Highly repeated. Daily reports, submittal logs, and pay application line items all recur on every project.

### Data & ML operations

Quality, governance, and readiness of the data and pipelines AI depends on.

### Tech stack & DevOps

How well the architecture and tooling support secure, scalable AI workloads.

### Talent & ways of working

The mix of skills, ownership, and processes that actually deliver AI outcomes.

### Governance & risk

The guardrails, controls, and oversight that make AI safe to scale.



**Fabio Ferraretto**

Chief Operating Officer



Fabio is a seasoned data and analytics leader with 20+ years of experience helping companies solve complex problems and improve performance. He holds a Master's in Logistics Systems Engineering and a BA in Construction Engineering from the University of São Paulo, plus a Lead Analytics & AI certification from Northwestern. He spent most of his career at Accenture Latin America, where he helped build the Analytics Group and led teams across Advanced Analytics and Supply Chain. He is one of the founders of Dhauz, a Quantum Rise company, a Data and AI solutions company based in Brazil. Since February 2026, Fabio is the COO of Quantum Rise.



**Josh Kohn**

SVP, Professional Services



Josh Kohn is the Senior Vice President of Strategy & Operations at Quantum Rise. Prior to joining Quantum Rise, Josh held various leadership roles at PathAI, including Associate Director of Program Management, where he was instrumental in managing cross-functional teams and timelines to improve patient outcomes through AI-powered pathology. He also brings experience from his time at ProPharma Group and Ventana Medical Systems, where he excelled in project management and strategic process improvements. Josh is PMP certified and holds an MBA from the Thunderbird School of Global Management.

---

**Hugh Seaton**

# Briefing: agents that work

*New Developments Make Developing and Using AI Agents Much Easier*



**Over the past two years,** the promise of AI agents outpaced the reality. Custom-built systems required heavy investment, produced mixed results, and left most organizations skeptical. One MIT study suggesting a 95 percent failure rate captured a real sentiment, even among firms that didn't take the number literally: agents weren't ready.

That changed at the start of 2026. A convergence of infrastructure maturity, developer experimentation, and new product releases moved agents from being laboratory curiosities to operational tools. Software developers, already using coding agents like

Claude Code for their day jobs, began applying those same tools to general business tasks: scheduling, email management, document workflows. Within weeks, Anthropic released Claude Cowork, an agentic platform designed for non-technical users. OpenAI followed almost immediately with an acquisition signaling the same strategic bet. The agentic era is no longer the future. It has arrived and it will define your company's performance. Perhaps surprisingly, however, agents' power lies not so much in their ability to automate tasks but to boost the worth of an organization's knowledge.

## What's new

The difference between 2025's failed pilots and 2026's working systems comes down to infrastructure, not AI model improvements. Three developments matter:

**Standardized connectivity.** AI models can now connect to external tools, enterprise software, databases, web services, through standardized protocols rather than bespoke integrations. This dramatically reduces the cost and fragility of connecting an agent to the systems where real work happens.

**Codified process knowledge.** Agents can be equipped with structured instructions, call them skills, recipes, or playbooks, that encode how to accomplish specific tasks within specific constraints. This is the missing link from 2025: agents failed not because the models lacked intelligence, but because they lacked context about how your organization works.

**Accessible harnesses.** An agentic harness is the system that orchestrates an agent's activities: assigning tasks, managing memory, routing to the right tools, and giving humans a control surface. Until recently, building a harness required engineering talent. Now, commercial platforms provide this out of the box.

Together, these developments mean that deploying an agent no longer requires a six-month pilot with a dedicated engineering team. A capable business user with clear process knowledge can stand up a working agent in days.

## Implications for Leaders

This new capability frees organizations to develop their people as owners of agentic processes, not just adopters of a pre-built solution. Ultimately the power of AI to transform organizations will be driven not by the state of the technology, but by the ability of those organizations to reimagine processes, integrate agents and drive radical improvements in speed, scope and customer value. Harnesses are the training ground for those capabilities.



### Hugh Seaton

Managing Director, Construction



Hugh is the Managing Director for the Construction Vertical at Quantum Rise. He started his construction technology innovation journey translating specifications into Chinese in Taiwan in the 1990s, and has been GM at the Construction Specifications Institute, founder of the AEC Hackathon, CEO of a construction software company, and now is leading the construction vertical at Quantum Rise, to help this essential industry transform with AI.

## Appendix:

## Examples of AI use case evolution in B2C Marketing

| AI AUTOMMATION                                                 |                            |                                                                   |                                                                         |                                                                               |
|----------------------------------------------------------------|----------------------------|-------------------------------------------------------------------|-------------------------------------------------------------------------|-------------------------------------------------------------------------------|
|                                                                |                            | ASSIST: HUMANS IN THE LOOP                                        | SCALE OUTPUTS: HUMANS ON THE LOOP                                       | AGENTIC EXECUTION: HUMAN OVERSIGHT                                            |
| Understand + Decide: insights, strategy, planning              | #1 Budget Planning         | AI drafts allocation scenarios; human decides                     | AI optimizes allocation vs. objectives; human approves                  | Dynamic reallocation executes in real time within approved guardrails         |
|                                                                | #2 VoC Insight Mining      | AI surfaces themes from reviews/social; human interprets          | AI clusters issues by segment; human validates findings                 | Always-on agent monitors VoC + sales + media; auto-triggers actions           |
|                                                                | #3 Audience Development    | AI drafts personas from data; human refines                       | AI builds/tests audiences from unstructured data; human approves        | Synthetic personas auto-evolve based on live campaign response                |
|                                                                | #4 Competitive Intel       | AI aggregates competitor activity; human synthesizes              | AI tracks launches, pricing, campaigns; human reviews intel             | Agent monitors competitors; recommends or initiates counter-actions           |
| Create + Orchestrate: content + personalized journeys          | #1 Personalized Journeys   | AI suggests segments and offers; human configures journeys        | AI runs recommendation models at scale; human monitors performance      | Real-time personalization adapts offers based on live behavior and context    |
|                                                                | #2 Content Generation      | AI drafts copy and creative variants; human edits and approves    | AI produces scaled variants across channels; human samples and approves | Continuous multivariate creative optimization deployed autonomously           |
|                                                                | #3 Service-to-Marketing:   | AI summarizes support signals; human creates retention campaigns  | AI maps support themes to offer recommendations; human approves         | Agent converts service interactions to retention/upsell actions automatically |
|                                                                | #4 PDP Enrichment          | AI drafts product copy and tags; human reviews each listing       | AI enriches catalog at scale; human QA spot-checks                      | Auto-optimization of PDP content based on live conversion signals             |
| Activate + Prove: media, measure + predict, optimization       | #1 Media Planning & Buying | AI suggests media plans and channel mix; human approves           | AI optimizes placements, pacing, creative rotation; human supervises    | Dynamic buying executes in real time to hit objectives within guardrails      |
|                                                                | #2 Campaign Performance    | AI generates dashboards and reports; human interprets and acts    | AI flags anomalies and surfaces insights proactively; human validates   | Autonomous agent adjusts campaign parameters; reports actions taken           |
|                                                                | #3 Experimentation         | AI suggests test designs; human sets up and runs                  | AI automates experiment setup and tracking; human oversees              | Continuous incrementality testing with automated readouts and actions         |
|                                                                | #4 Influencer Activation   | AI identifies and ranks creator candidates; human selects         | AI manages outreach and performance tracking; human monitors            | End-to-end sourcing, vetting, contracting, and optimization automated         |
| Governance + Enablement: compliance, data quality, martech ops | #1 Data Quality            | AI identifies pipeline anomalies; human investigates and resolves | AI monitors data flows continuously; human oversees resolution          | Auto-remediation of data issues; opens tickets and alerts on marketing impact |
|                                                                | #2 Martech Ops             | AI assists with setup, troubleshooting, and playbook creation     | AI automates routine workflows; human oversees exceptions               | Cross-functional orchestration with autonomous routing and execution          |
|                                                                | #3 Legal & Privacy         | AI highlights consent and privacy risks; human makes decisions    | AI enforces consent and privacy rules across campaigns; human monitors  | Real-time compliance enforcement embedded in execution workflows              |
|                                                                | #4 Brand Compliance        | AI flags brand and voice issues in content; human reviews         | AI enforces brand rules across assets at scale; human audits            | Blocks non-compliant content pre-publish; routes edge cases; logs decisions   |



  **in**

 **QuantumRise**