

OCTANT ADVISORY

PRIVATE EQUITY WHITE PAPER

AI Performance Diligence

How private equity finds the organizational risks that keep AI from reaching the P&L, before the deal and across the hold

A practical framework for investment teams and operating partners

JULY 2026

octantadvisory.com

EXECUTIVE SUMMARY

AI spending is climbing fast. Few companies can show it in the P&L.

U.S. private firms are on track to spend about \$280 billion on AI in 2026, and per-employee spending rose roughly 50% in a single year (Federal Reserve Bank of Atlanta). The capital is committed. Whether it reaches the P&L is a separate question.

COMPOSITE ILLUSTRATION

A portfolio company had six AI pilots, capable technical leaders, and a growing software budget. What it did not have was an executive accountable for the financial result, a shared baseline, or a plan to convert saved time into earnings. The technology worked. The value-creation system did not.

This is the pattern we keep finding. AI is present in most companies. That presence does not make it an asset. Whether it becomes one depends on the organization built around it.

INVESTMENT IMPLICATION An AI initiative is not an asset simply because it exists. Its value depends on whether the company can govern, scale, measure, and redesign the business around it.

For private equity, this creates a distinctive problem. The hold period imposes a clock. Fragmented systems, thin management capacity, unclear accountability, and weak data disciplines can consume that clock before a promising pilot becomes a repeatable earnings lever.

Traditional technology diligence can establish what systems, data, vendors, and models a company has. It often does not answer the more consequential question: can this management team and operating model convert those assets into measurable performance within the investment horizon?

A better underwriting question

The relevant question is not “Does this company use AI?” It is:

What organizational conditions must be true for AI to affect revenue, margin, working capital, risk, or exit multiple, and which of those conditions are currently missing?

Octant Advisory calls this AI performance diligence: an independent, evidence-based assessment of the organizational, leadership, and technology conditions that determine whether AI investment can perform.

Three conclusions for PE leaders

- Treat AI readiness as an underwriting variable, not a post-close technology workstream.
- Sequence remediation around binding constraints; do not fund a backlog of disconnected use cases.
- Measure value against the deal thesis, with named owners, baselines, milestones, and decision rights.

01 | THE INVESTMENT CONTEXT

AI has entered the value-creation plan but not yet the earnings bridge

PE firms face pressure to find operational value in a market where leverage is less forgiving and hold periods are longer. AI appears to offer a new path: better pricing, faster sales cycles, improved service productivity, leaner back offices, stronger forecasting, and differentiated products.

The opportunity is real. Turning it into earnings takes deliberate work.

ADOPTION	Roughly a third of larger firms (250+ employees) now use AI, and the share is climbing (U.S. Census Bureau, 2026).
SCALING	Most deployments stall as isolated pilots and never reach production at scale.
FINANCIAL IMPACT	Few companies can trace the result to the P&L or to cash flow.
HIGH PERFORMANCE	A small minority build the operating system that turns AI into durable earnings.

The gap is organizational. Technology is rarely the binding constraint.

The pattern holds across the companies we work with. The divide is organizational more than technical. It comes down to how work is redesigned and how deeply new tools are embedded in real workflows. That is where AI value is won or lost.

THE PE TRANSLATION Buying more technology is often the most visible action, and the least sufficient response.

Why portfolio companies are particularly exposed

- Management bandwidth is scarce, and AI competes with core transformation priorities.
- Data and process debt are frequently inherited from years of pragmatic underinvestment.
- A small number of leaders may hold critical institutional knowledge, creating key-person risk.
- Use cases may be selected for novelty or ease, not their contribution to the investment thesis.
- Benefits may appear as “hours saved” without a mechanism to convert capacity into margin or growth.

02 | THE PERFORMANCE GAP

The hidden risk sits between the pilot and the P&L

An AI pilot can be technically successful and economically irrelevant. It can generate accurate outputs, attract users, and still fail to change a workflow, a decision, a customer experience, or a cost structure.

That gap is created by dependencies outside the model. A pricing engine cannot improve margin when commercial leaders do not trust it. A service copilot cannot reduce cost when workforce plans do not capture released capacity. A forecasting tool cannot improve working capital when business units use different definitions and data.

CORE PRINCIPLE Technology enables AI value. The organization is what captures it.

Five failure patterns PE teams should recognize

The orphaned mandate	AI is encouraged, but no executive owns an enterprise outcome or has authority to change workflows.
The use-case portfolio	Many pilots proceed, but none is large enough, or governed tightly enough, to move the value-creation plan.
The technology shortcut	A vendor is selected before process, data, risk, adoption, and economics are defined.
The unpriced dependency	Critical data cleanup, integration, governance, or talent work is absent from the business case.
The unmeasured benefit	Activity metrics substitute for financial baselines, accountable owners, and verified value capture.

These are not reasons to slow AI indiscriminately. They are reasons to diagnose the system around AI before allocating additional capital. The objective is speed with sequence: remove the few constraints that prevent everything else from moving.

03 | THE DILIGENCE FRAMEWORK

Eight dimensions determine whether AI investment can perform

These eight dimensions make up the AIR-O, the organizational lens of Octant's AI Reality Index™. Octant advisors score each dimension on a five-level scale from evidence, not from how the team rates itself. The output is not a maturity label. It is the two or three binding constraints, in the order they must be addressed, that determine what the company should fund, change, stop, govern, or hire.

	DIMENSION	DILIGENCE QUESTION
1	Strategic Mandate	Is AI tied to a funded enterprise ambition and the deal thesis?
2	Executive Alignment	Do leaders agree on outcomes, priorities, ownership, and tradeoffs?
3	Governance & Decision Rights	Are AI decision rights clear, accountability defined, and risk actively managed?
4	Data & Technology Strategy	Is data governed as an enterprise asset, with clear build, buy, and partner decisions?
5	Operating Model & Structure	Can business and technical teams move from use case to changed workflow?

	DIMENSION	DILIGENCE QUESTION
6	Leadership Capability	Can accountable leaders carry the transformation under pressure?
7	Workforce & Culture Readiness	Will roles, incentives, skills, and trust support adoption?
8	Value Measurement & Improvement	Are baselines, KPIs, benefits, and reinvestment decisions explicit?

The score is not the decision. The pattern across the dimensions is. Low value measurement may make every use case speculative. Weak executive alignment may prevent a shared portfolio of priorities. A thin leadership bench may make an otherwise credible plan impossible within the hold period.

WHAT MAKES THE AIR READ DIFFERENT

Two scores, not one. Each dimension gets a Claimed score, what leaders report, and an Evidenced score, what artifacts and observed decisions support. The Evidenced score is the one that counts. The distance between them is the Reality Gap, and a wide gap is often the most important finding: the organization believes work is done that the evidence cannot confirm.

Disagreement is evidence. When two executives score the same dimension three or more points apart, that gap, not the average, is frequently the primary finding.

Constraints, not weaknesses. We isolate the two or three lowest dimensions blocking everything else, in the order they must be addressed, rather than cataloguing every gap.

Scored from evidence. Advisors assign the scores from interviews, document review, and a station-by-station trace of the company's own AI initiatives, not from a self-rating.

Graded in six months. Every recommendation ships as a measurement contract with a baseline, target, named owner, leading indicator, and test date, then is re-scored at month six against what we predicted.

04 | ACROSS THE INVESTMENT LIFECYCLE

Use AI performance diligence at three decision points

PRE-ACQUISITION

Underwrite capability and exposure

- Test whether claimed AI advantages are repeatable, governed, and economically attributable.
- Identify hidden dependencies in data, integration, leadership, and process redesign.
- Estimate the capital, talent, and time required to reach the value-creation case.

FIRST 100 DAYS

Convert the thesis into a sequenced agenda

- Name the executive owner and establish board-level decision rights.
- Select a small number of value pools tied directly to the investment thesis.
- Baseline economics before deployment and remove the first binding constraints.

HOLD & EXIT**Scale proof and strengthen the equity story**

- Track adoption, operational outcomes, and verified financial value separately.
- Standardize reusable capabilities without forcing identical use cases across the portfolio.
- Document durable governance, operating routines, and evidence of repeatable performance.

The evidence standard should rise with the decision

Early screening can rely on targeted interviews and document review. A material underwriting claim requires stronger evidence: named owners, workflow maps, data lineage, vendor obligations, controls, adoption data, unit economics, and management's demonstrated ability to execute.

By exit, the buyer should not have to accept "AI-enabled" as a narrative assertion. The company should be able to show what changed, what value resulted, what controls endure, and why the capability is transferable beyond a few individuals.

05 | THE FIRST 100 DAYS**Sequence the work around constraints, not enthusiasm**

The first 100 days should establish the management system for AI value. It should not become a competition to launch the most pilots.

1. Anchor

Translate the deal thesis into two or three enterprise outcomes. Define the financial baseline and the operating metric that AI must move.

2. Diagnose

Score the eight AIR-O dimensions from executive interviews, document review, and a Value Path Trace of the company's own initiatives, assigned by advisors against a defined rubric. Treat wide disagreement between executives as evidence.

3. Prioritize

Select use cases by value pool, feasibility, time to proof, dependency load, and reversibility, not executive sponsorship alone.

4. Govern

Set an accountable executive, decision rights, risk thresholds, funding gates, and a cadence that connects delivery to the operating review.

5. Redesign

Map the end-to-end workflow. Specify how roles, decisions, controls, incentives, and capacity will change when the tool is deployed.

6. Prove

Track adoption, operational effect, and financial capture separately. Stop or redesign initiatives that cannot cross those gates.

A value gate that prevents false positives**GATE****QUESTION****EVIDENCE**

Adoption	Did the intended users change behavior?	Active use in the target workflow
Operational	Did the process outcome improve?	Cycle time, conversion, quality, yield
Financial	Was value captured in the P&L or cash flow?	Verified revenue, cost, margin, or working capital

06 | THE PORTFOLIO PLAYBOOK

Five interventions for operating partners

01	Create a common evidence standard	Use one diagnostic language across the portfolio while allowing priorities to reflect each company's thesis, sector, and starting point.
02	Fund foundations with explicit economics	Data, integration, security, and governance are not overhead when they are prerequisites to an identified value pool. Price them into the case.
03	Match leadership to the transformation	Determine whether the company needs a builder, transformer, or operator, and whether the current executive has the authority and capability to deliver.
04	Separate reusable capability from local use cases	Standardize controls, vendor disciplines, measurement, and talent where scale helps. Keep commercial and operational use cases close to the business.
05	Turn exit claims into auditable proof	Maintain a value ledger linking each initiative to its baseline, adoption, operating effect, financial capture, owner, and supporting evidence.

What the board should see

A concise quarterly view should show: the value pools being pursued; the constraints that threaten them; the owner of each intervention; adoption, operational, and financial indicators; risk exceptions; and the next capital-allocation decision.

DISCIPLINE A roadmap is valuable only if it changes a decision: what to fund, stop, sequence, govern, or hire.

07 | FIELD CHECKLIST

Twelve questions before approving the next AI dollar

LENS	QUESTION
☐ Thesis	Which deal-thesis outcome will this initiative move?
☐ Baseline	What is the current financial and operational baseline?
☐ Owner	Which executive owns the outcome, not merely the implementation?
☐ Workflow	What decision or end-to-end workflow will change?
☐ Dependency	Which data, process, integration, or talent constraints must be removed first?
☐ Economics	Are all enabling costs included in the return case?
☐ Adoption	What behavior must change, and how will that behavior be observed?
☐ Capacity	How will saved time convert into margin, throughput, quality, or growth?
☐ Governance	Who can approve, pause, escalate, or retire the initiative?
☐ Risk	What failure modes could create customer, regulatory, IP, security, or reputation exposure?
☐ Proof	What evidence will establish operational and financial impact?
☐ Exit	Would a buyer view the capability as durable, transferable, and independently verifiable?

A “no” does not automatically disqualify an initiative. It identifies an underwriting assumption or execution dependency that should be made explicit, owned, costed, and sequenced.

08 | CONCLUSION

The firms that win will underwrite the organization around the technology

AI will become less differentiating as access to models and tools broadens. The durable advantage will come from the ability to choose the right value pools, redesign work, govern risk, align leaders, build trust, and verify financial impact faster than competitors.

For private equity, that capability belongs inside the investment process. It should shape diligence, the value-creation plan, leadership decisions, portfolio governance, and the exit narrative.

The practical shift:

Read AI readiness as evidence about whether the organization can convert investment into performance within the hold period, rather than as a technology maturity score.

How Octant Advisory helps

Octant Advisory is an independent AI diagnostics and advisory firm. We help CEOs, boards, investors, and operating partners identify the organizational conditions constraining AI performance and sequence the changes required to correct them.

The AI Reality Index™ (AIR) combines three connected lenses:

- AIR-O: organizational readiness across strategy, governance, operating model, workforce, and value measurement.
- AIR-T: the technical foundation, opened when the Data & Technology Strategy dimension points to a technical root cause.
- AIR-L: the leadership traits that predict transformation success, opened when the Leadership Capability dimension scores below threshold.

Octant does not sell technology or implementation. Independence is the point: the recommendation follows the evidence, not a platform sale or downstream delivery contract.

BEGIN WITH THE EVIDENCE Request a confidential 30-minute briefing to review the five organizational risks most likely to impair AI value in a target or portfolio company.

octantadvisory.com | info@octantadvisory.com | (904) 638-0066

SOURCES & METHODOLOGY

Selected sources

Forbes Research. *AI's ROI Reality Check: Are Companies Measuring What Matters.* October 8, 2025.
<https://www.forbes.com/sites/forbes-research/2025/10/08/ai-roi-measurement-challenges-forbes-survey-2025/>

Federal Reserve Bank of Atlanta. *How Much Are Firms Spending on AI (and What Will Happen to Headcounts)?* May 2026.
<https://www.atlantafed.org/research-and-data/publications/policy-hub-macroblog/2026/05/06/how-much-firms-spending-on-ai-and-what-will-happen-to-headcounts>

U.S. Census Bureau. *Business Trends and Outlook Survey: Artificial Intelligence Use in Businesses.* 2026.
<https://www.census.gov/library/stories/2026/05/ai-use-businesses.html>

Forbes. *Lack of ROI From AI Reveals Pitfalls for Business.* February 5, 2026.
<https://www.forbes.com/sites/johnwerner/2026/02/05/lack-of-roi-from-ai-reveals-pitfalls-for-business/>

Octant Advisory. *The AI Reality (AIR) Index and Independent AI Diagnostics.* accessed July 2026.
<https://octantadvisory.com/approach.html>

Methodology note

This paper synthesizes independent, third-party research with Octant Advisory's practitioner framework. The AIR-O dimensions surface diagnostic hypotheses to test with the leadership team, not a statistical prediction of investment returns. Findings for any company should be based on interviews, documentary evidence, operating data, and a defined rubric rather than management self-assessment alone.

About Octant Advisory

Octant Advisory is an independent advisory firm helping CEOs, boards, senior operators, and investors convert AI investment into measurable enterprise performance. Its work focuses on the organizational, leadership, and technology conditions that determine whether AI can scale.

© 2026 Octant Advisory LLC. All rights reserved. This paper is for informational purposes and does not constitute investment, legal, accounting, or technology advice.