

The PE playbook: translating AI risk into deal structure

JULY 16, 2026 4 MIN READ



Related Expertise

- [Artificial Intelligence](#)
- [Mergers and Acquisitions](#)
- [Private Equity](#)
- [Technology](#)

Authors: [Brian Gray](#), [Sam Ip](#)

Key Takeaways

- Sponsors should assess broader operational assumptions beyond traditional software or technology diligence, such as rights to training data, dependencies on third-party providers and other legal/technical foundations for supporting future growth.
- Traditional transaction tools, including representations and warranties insurance, are not designed to cover AI risks that can emerge after closing.
- Protecting the investment thesis requires more than legal compliance, and sponsors should ensure the target has the rights, governance and key talent necessary to execute its post-close strategy.

When PE sponsors underwrite an AI-driven platform or a target leveraging proprietary models, the technical blind spots identified in [our March 10 article](#) — such as dynamic data dependencies and unpredictable future performance — cannot simply be outsourced to standard legal diligence. They directly impact the investment thesis, purchase price, transaction structure and, ultimately, exit value.

Rather than treating AI as another software acquisition, sponsors should adapt traditional transaction tools across the various stages of investment lifecycle, from underwriting and diligence to transaction structuring and executing on the investment thesis post-close. The legal framework should reinforce the business assumptions underpinning the investment.

Looking beyond traditional software diligence

Traditional technology diligence asks whether the target owns its IP, complies with applicable laws and has assignable customer contracts. AI businesses require sponsors to underwrite a broader set of operational assumptions that may directly impact valuation.

In addition to traditional legal diligence, sponsors should assess whether the target has durable rights to its training data, the extent to which it depends on third-party foundation model providers, whether compute commitments and infrastructure costs align with the business model, the maturity of its AI governance framework, the ability to scale customer data rights in support of the investment thesis and whether ongoing model performance

depends disproportionately on a small number of key technical personnel.

Structuring the transaction and adapting traditional risk allocation tools

Representations and warranties insurance (RWI) remains ill-equipped to cover forward-looking AI risks, such as algorithmic drift or future intellectual property claims arising from evolving training sets or changes in third-party foundation model providers. More fundamentally, traditional representations and warranties are designed to address historical facts existing at signing or closing. Adapting AI systems can create risks that materialize after closing.

In Canada, the legal framework governing AI use of data continues to evolve. Proposed federal privacy reforms under Bill C-36, together with Quebec's Law 25, increase expectations around transparency and automated decision making. Taken together, these developments do not create a standalone liability regime, but heighten regulatory exposure associated with AI and AI-related data assets.

Because these compliance gaps are not included in standard RWI policies, deal teams may need to consider targeted indemnities, bespoke covenants or arrangements, or even post-closing commitments around data, AI governance or regulatory compliance.

Many of the most significant risks arise not from historical legal compliance, but from the continued ability of AI systems to perform, evolve and create value after closing.

Sponsors also may utilize earn-outs to bridge valuation gaps on high-growth AI targets, but traditional financial metrics (e.g., EBITDA or revenue multiples) may fail to capture the true health of an algorithmic enterprise. Sponsors may begin to start thinking about earn-outs tied directly to AI system performance, workflow adoption rates and model stability.

However, an earn-out tied to algorithmic output will likely come with litigation risks. Sponsors must ensure clear drafting regarding what constitutes a "qualifying output," who controls the model's training environment post-close and how disputes are resolved if model behaviour shifts or degrades due to buyer-driven portfolio integration decisions rather than seller-side deficiencies.

Protecting investment thesis post-close

A recurring issue in practice is that a sponsor discovers post-closing that it lacks the legal infrastructure necessary to execute its deal thesis. The target may have lawfully deployed its AI solution historically, but lacks the rights necessary to expand into new geographies, integrate newly acquired businesses, leverage customer datasets across portfolio companies or deploy the technology in new commercial settings.

Diligence must map the intended future use against legacy customer consents and data licenses, rather than simply verifying whether historical use was lawful. Where the investment thesis depends on cross-portfolio data integration, closing conditions or post-closing covenants may be required to obtain additional rights.

Because the value of AI enterprises often resides as much in the people responsible for developing and maintaining the AI systems as in the underlying code itself, traditional management retention arrangements should be supplemented with incentives for key AI talent (e.g., data scientists and ML engineers), which may be incredibly scarce and in high demand. Depending on the investment thesis, sponsors may also require governance tied to

preserving institutional knowledge, ongoing model continuity and long-term accountability for regulatory compliance.

Conclusion

AI transactions require sponsors to think beyond traditional software acquisition playbooks. Many of the most significant risks arise not from historical legal compliance, but from the continued ability of AI systems to perform, evolve and create value after closing.

For sponsors, the objective is therefore not simply identifying AI risk, but translating that risk into underwriting assumptions, transaction structure and post-closing governance that preserve the investment thesis through to exit.