

AI transactions: why traditional approaches fall short and how to adapt

MARCH 10, 2026 9 MIN READ



Related Expertise

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

Authors: [Sam Ip](#), [Matthew T. Oliver](#), [Calvin Leung](#), [Jordan Geist](#)

Key Takeaways

- Transaction structures and terms are evolving, focusing on AI-specific diligence, longer warranty periods and adaptable structures to manage AI risks effectively.
- Post-closing risks can arise from deploying AI in new contexts, leading to potential legal and regulatory issues not identified during due diligence.
- Buyers should prioritize securing data rights and control over AI assets, as ownership alone doesn't fully ensure the ability to adapt AI systems post-acquisition.

Over the past year, artificial intelligence has moved from a diligence consideration to a central driver of deal thesis and valuation in technology transactions. This shift is not limited to acquisitions of AI-native companies. Increasingly, AI is embedded across products, services and internal workflows, even where the target is not being marketed as an AI business. Market data reflects this trend, with a growing proportion of technology deals now involving an AI component that is often embedded in the core operations and value proposition of the deal. As a result, the heart of what is being bought and sold in technology transactions is shifting.

Building on Osler's prior discussions of AI-focused [diligence](#) and [representations and warranties](#), this Update explores recurring questions we have been discussing with clients this past year, framing the reasons why conventional deal mechanics often fall short of surfacing the most material AI risks and how transaction practice is adapting.

1. Why do traditional transaction frameworks struggle when AI is central to value?

Traditional transaction frameworks are designed to assess largely static assets and historical compliance, assuming risk can be identified through diligence, confirmations of ownership and point-in-time representations. The nature of AI systems strains this model.

Diligence in technology transactions has traditionally focused on IP ownership, license validity and lawful data collection. However, the assets, people and business supply chain underlying AI systems are often more complex and dynamic. Confirmation of ownership alone is insufficient as AI systems often rely on licensed data, third-party models, open-

source components and evolving deployment permissions, as well as specialized personnel beyond traditional software developers, including data scientists and those responsible for ongoing monitoring and human oversight.

Similarly, AI system output is probabilistic and context-dependent. Some of the more noteworthy AI risks can arise from outputs: discriminatory or biased decisions, unsafe or inaccurate recommendations, or outcomes that fail to meet sector-specific regulatory standards. These risks frequently surface only once systems are deployed, scaled or integrated into new contexts post-closing, as the behavior of an AI system can change when systems are deployed into new contexts. Traditional representations and warranties are calibrated to assess historical noncompliance, not future outputs, and may only provide limited assurance with respect to future system performance or post-closing use.

To address the shortfalls of traditional transaction frameworks, buyers do not solely rely on legal diligence. AI systems are increasingly evaluated in controlled environments, allowing buyers to assess system behaviour and performance under controlled conditions. In effect, this is a “try before you buy” dynamic, where buyers seek validation of AI system behaviour before closing.

2. What AI risks typically emerge after closing and how should deal teams anticipate them?

AI risks can crystallize post-closing even if nothing was “wrong” at closing. This is because AI systems are often deployed post-closing in broader or different contexts — for instance, with new customers, in new geographies or in different decision-making functions. In many transactions, AI systems were developed and validated against the target’s historical use case, data and operating environment, though the buyer’s deal thesis may contemplate materially different applications post-close. These changes to the application of the model can create legal, regulatory or use case-specific risks that were not identified at the time of the deal.

A recurring issue in practice is that buyers discover post-closing that they lack the data rights necessary to implement the AI system in a manner consistent with the deal thesis. This can arise where data licenses, consents (particularly in the case where personal information is involved) or collection practices permit the target’s use of the AI system, but do not extend to buyer’s intended post-closing use cases, customers or geographies. In some cases, personal information consents are too narrow or, in other cases, licensed data sets do not permit retraining, expansion or applications to new scenarios. This is particularly relevant where key datasets are licensed and subject to use restrictions. These types of data license restrictions can be anticipated through appropriate diligence, including asking questions about the data, its provenance and use rights throughout the transaction.

Deal teams should consider referencing evolving AI standards, such as commitments to comply with leading frameworks (e.g., NIST AI RMF or ISO/IEC 42001), as part of diligence to assess whether the target has a consistent approach to AI governance, model monitoring and human oversight. Where gaps are identified in diligence, these frameworks can be used as a basis for targeted representations or post-closing remediation, mitigating the risk that AI models may become inaccurate or biased and aligning with industry best practices.

3. In an AI-driven transaction, what AI assets should the buyer consider needing to control?

In AI transactions, the scope of rights over data, models and expertise often matters more than formal ownership. A buyer therefore risks acquiring ownership of a target without securing the necessary rights to scale or adapt AI systems in a manner consistent with their deal thesis. Key areas of rights and control include the following:

- **rights over data access and future data flows.** Datasets are often licensed rather than owned, with restrictions on retention, retraining, geographic expansion or new use cases. If the buyer is looking to extend the use of AI system to new products, customers or jurisdictions, this should be contemplated when evaluating the scope and transferability of the license.
- **rights over third-party dependencies.** AI systems often rely on various external components, including APIs, pre-trained models or open-source libraries. The buyer must be able to secure appropriate rights to continue using, modifying and scaling these dependencies post-close, and have visibility into any obligations triggered by scaling or modifying the AI system.
- **retention and availability of key personnel and expertise.** AI systems frequently depend on specialized personnel whose expertise is not fully captured in documentation or code. Furthermore, evolving regulatory and best-practice standards view human oversight as integral to responsible AI use. The buyer should control key technical talent and ensure the transfer of operational know-how, as the loss of key personnel can not only impair system performance but also create compliance issues.

Given these factors, post-close AI risk is increasingly about whether the buyer has secured the rights, expertise and operational control necessary to operate and scale the AI system after closing.

4. How are transaction structure and terms evolving to manage AI risk?

As discussed, AI risk often emerges through post-closing use, scaling and modification of AI systems. As a result, both transaction structure and deal terms are evolving to address these unique risks.

On the structural side, deals are shifting from instantaneous exits toward time-based risk allocation, as many forward-looking AI risks cannot be fully addressed through traditional transaction structures. Common structural approaches include

- acqui-hire and hybrid deals where the core value is concentrated in talent
- earn-outs tied to usage, workflow adoption, output quality or margin stability rather than revenue
- staged partnerships or licensing arrangements used as validation steps before a full acquisition
- asset deals or carve-outs to isolate model and data pipelines while avoiding legacy liabilities

The choice of the transaction structure materially influences post-closing AI risk and complexity. In a share sale, the buyer typically acquires the target's existing data rights, licenses and contractual arrangements as a going concern. This may simplify continuity but can also mean inheriting legacy restrictions or consent limitations that affect how AI systems may be used or scaled post-closing. By contrast, asset sales and carve-out transactions often require explicit assignment or renegotiation of data licenses, API access agreements and third-party model dependencies, and may trigger change-of-control provisions or consent requirements that could limit the buyer's ability to continue operating AI systems in their current form. These transactions also frequently necessitate transition services agreements to address shared data, infrastructure or personnel that remain with the target but are operationally critical to the AI system being acquired.

Given that AI systems often depend on ongoing data flows, specialized talent and external dependencies that may not transfer seamlessly, deal teams should carefully evaluate how the chosen transaction structure affects the scope and duration of transition services, as well as the risk that shared data environments or personnel arrangements may impair the buyer's ability to deploy, monitor or retrain AI systems in a manner consistent with the deal thesis.

Beyond structure, transaction terms are increasingly including AI-specific diligence questions and representations and warranties. These representations and warranties, often with longer survival periods and higher indemnity caps, can help manage AI risk post-close and may address the rights to use data, permitted uses of that data and the absence of undisclosed third-party dependencies or restrictions that would limit post-closing deployment or scaling of the AI system. The market has seen a tendency toward expansive, "kitchen-sink" AI representations and diligence that attempt to address every conceivable AI issue. In our experience, however, this approach often adds cost and friction without meaningfully improving risk mitigation. Use case-driven tailoring is critical to avoid unnecessary delay and cost while still addressing AI risks that are more likely to be material to value.

Finally, parties should be aware that representation and warranty insurance is often a poor fit for underwriting forward-looking AI usage risks. While RWI may cover losses arising from breaches of well-drafted factual representations, it generally does not insure post-closing risks such as model performance, drift or buyer-driven deployment choices.

5. What does a practical AI transaction playbook look like for boards and deal teams?

A practical playbook for buyers begins with the deal thesis and determining whether and to what extent AI is material to the transaction's value. That assessment informs whether the transaction should be approached as an AI-driven deal at all, before delving into detailed diligence or transaction structure and terms. In some cases, that may involve an early, limited-scope diligence to assess whether foundational elements are viable (for example, whether the company has the necessary data rights to support the intended use case) before committing to full diligence.

For sellers, their board and investors should consider AI risk as an exit issue, not just an operating issue. Boards and deal teams should align early on acceptable earn-out exposure, founder and management team retention expectations and ongoing governance commitments, and prepare their investors for deferred value realization and non-traditional exit mechanics.

Where AI is central to transaction value, boards should expect AI considerations to move to the forefront of deal approval and ensure that appropriate technical, legal and operational expertise informs decision-making. Boards and deal teams should also plan for post-closing oversight and governance as part of the overall transaction, rather than as an afterthought. Ultimately, the goal for both boards and deal teams is not to eliminate all risk in an AI transaction, but to identify, allocate and manage AI-related risk that is proportional to the role AI plays in the deal thesis and the intended scale and deployment of AI post-close throughout the AI lifecycle.