Addressing AI Risks in Merger and Acquisition Transactions

By Martin Florman, Daniel S. Dubyak | July 29, 2026

RELATED PRACTICES AI and Technology, Corporate Governance, Corporate, Data Centers and Digital Infrastructure, Mergers and Acquisitions

Artificial intelligence has moved from a niche technology concern to a mainstream operational tool embedded across virtually every industry. Generative AI platforms, machine learning models, automated decision systems and vendor-hosted AI tools are now used routinely by businesses of all sizes, often without a complete inventory of where and how AI operates within the organization. Many companies have adopted AI tools faster than their governance frameworks, policies or contractual protections have kept pace. At the same time, the legal and regulatory landscape surrounding AI is evolving rapidly. Federal and state legislatures, the FTC, the SEC and sector-specific regulators are continuously issuing new guidance, enforcement actions and proposed rules addressing AI transparency, bias, data use and accountability.

For buyers and sellers in M&A transactions, this means that diligence and transaction drafting must be dynamic, as what was adequate even six months ago may no longer be sufficient to allocate AI-related risks.

1. **How AI Is Used in Target Businesses:** AI is no longer limited to a company's IT department or R&D department. Today, AI is embedded in products and services sold to customers, internal business operations, third-party software and cloud vendor platforms, customer-facing automated tools (such as chatbots and recommendation engines) and employee use of public or enterprise generative AI tools. Businesses rely on AI for coding and product development, marketing content creation, HR screening and hiring, legal and compliance analysis, financial analytics, pricing and underwriting, procurement and operational decision-making. Because AI use is so pervasive, a target company in an M&A transaction may have significant AI exposure even if it does not consider itself an "AI company." Buyers should understand the full scope of a target's AI footprint, including AI tools used by employees without formal approval, data fed into third-party models and AI-generated outputs relied upon in business-critical functions.
2. **Risks Created by AI Use:** The breadth of AI adoption and use creates corresponding risks. Key AI-related risks in M&A transactions include (i) leakage of confidential information (e.g., employees inputting proprietary data into public AI tools), (ii) privacy and data protection violations, (iii) cybersecurity vulnerabilities and unauthorized model access, (iv) intellectual property ownership disputes and infringement claims arising from AI-generated content or training data, (v) inaccurate AI outputs and risk of reliance on inaccurate information, (vi) bias and discrimination in employment, credit, pricing or customer eligibility decisions, (vii) regulatory noncompliance with emerging AI laws, (viii) vendor lock-in and dependency on third-party AI providers, (ix) failure to disclose AI use to customers or regulators, (x) open-source or third-party license contamination, (xi) absence of audit trails for AI-driven decisions and (xii) employee misuse of AI tools. Additionally, in transactions involving transition services or shared systems, failure to properly allocate AI systems, underlying data and trained models between the buyer and the target company can create post-closing operational disruption, data commingling and

ongoing liability exposure. These risks should be identified, quantified and addressed across the full suite of transaction documents (i.e., not treated as a generic IT issue).

3. **Best Practices in Drafting Transaction Documents:** Addressing AI risk effectively requires a cross-document approach. In the early NDA phase, definitions of "Confidential Information" and permitted use restrictions should explicitly address AI tools, prohibiting the input of disclosed information into AI systems without prior consent. Diligence requests should be expanded to include a full AI inventory, training data sources, vendor agreements, internal AI policies, incident history, regulatory notices, IP claims, model governance protocols, human oversight mechanisms and customer-facing AI disclosures. In purchase agreements, AI-specific representations and warranties should address ownership of AI-generated IP, compliance with AI-related laws, data use in model training and the absence of undisclosed AI-related claims or incidents. Interim operating covenants should restrict material changes to AI systems, vendor relationships or data practices between signing and closing. Indemnification provisions, special escrows or purchase price adjustment mechanisms should be considered for known or contingent AI risks identified during diligence. Employment and consulting agreements should include clear assignment of AI-generated work product and acceptable use policies. Transition services agreements should address ownership of and access to AI systems, data segregation, service levels for AI-dependent processes and wind-down or migration rights. Other specific ancillary documents should include vendor-consent (and license-transfer analysis for AI-related contracts).

Practitioners should also reassess on an ongoing basis whether boilerplate AI-use representations continue to serve their intended purpose. Representations drafted even a year ago tended to start from the premise that AI use was a novel and material risk factor warranting disclosure. As tools such as generative search summaries, Microsoft Copilot (for routine drafting) and Claude or similar tools (for spreadsheet and data analysis) have become standard business tools across virtually every industry, the more probative diligence question in many deals may no longer be whether a target uses AI, but why it would not. A representation premised on the novelty of AI use risks becoming outdated as that premise itself becomes outdated. There is also a practical limitation inherent in disclosure-based AI representations: if use of a generative AI tool has already compromised trade secret protection or another confidentiality safeguard (for example, by inputting proprietary information into a public model), a representation that merely requires the target to disclose that use after the fact does little to remedy the harm, which, if it has occurred, has already occurred by the time the representation is given. Counsel should be candid with clients that a generic AI-use representation functions primarily as a disclosure and risk-allocation mechanism going forward, not as a substitute for the confidentiality and data-handling diligence that should have occurred earlier in the relationship.

Given this shift, AI-specific representations continue to add the most value in a narrower set of circumstances. The first is where AI use is independently subject to regulation, such as chatbot disclosure requirements that inform users they are interacting with an AI system or requirements governing the use of AI in employment decisions (e.g., hiring, promotion or termination) that carry bias-testing, notice or audit obligations. Even here, counsel should consider whether a generic compliance-with-laws representation already captures these regulated use cases, which would make a standalone AI representation duplicative rather than additive. The second circumstance, which is often more consequential, is where the target's core value proposition is its proprietary AI technology or related intellectual property. In those transactions, a generic AI-use representation is frequently the wrong tool. Buyers are better served by tailored representations and diligence focused on ownership of the AI models and underlying IP, the provenance and licensing of training data and the adequacy of technical and contractual protections around the model itself, rather than by broad disclosure obligations directed at incidental AI use across the business.

A related question has arisen with respect to the AI-use restrictions now regularly requested in M&A NDAs. An express prohibition on inputting confidential information into AI systems functions less as a durable, freestanding protection and more as an interim reminder that raises counterparty awareness of AI-related confidentiality risk while internal policies and practices are still catching up to the technology. As AI tool use continues to normalize across the workforce, and to the extent existing confidentiality and trade secret obligations are understood to already reach such conduct, the incremental value of a standalone AI-use restriction in the NDA may diminish over time, even though it remains a useful diligence marker and awareness-raising device in the near term.

4. **Sample Provision Language:** The following sample provisions (i) are illustrative, (ii) should be adapted to the specific facts, risks and structure of each transaction and (iii) can in no way be taken as legal advice. They merely demonstrate the type of language that experienced M&A counsel may incorporate/tailor to address AI-related risks and may become obsolete as the legal and regulatory landscape surrounding AI evolves. In particular, the purchase agreement representation below illustrates a broad, disclosure-based approach; as discussed above, counsel should consider narrowing or supplementing this language where the AI use at issue is independently regulated, already captured by a general compliance-with-laws representation or where the target's business is built around proprietary AI technology warranting more tailored, IP-focused representations and diligence.

   - **NDA — AI Restriction on Confidential Information**

     "*The Receiving Party shall not, and shall cause its Representatives not to, input, upload, or otherwise make available any Confidential Information of the Disclosing Party to any artificial intelligence system, machine learning model, large language model, or similar technology (whether hosted internally or by a third party) without the prior written consent of the Disclosing Party.*"

   - **Purchase Agreement — AI Representation and Warranty**

     "*The Company has made available to Buyer a complete and accurate inventory of all material artificial intelligence and machine learning systems used in or developed for the Business, including a description of (i) the purpose and function of each such system, (ii) the data sets used to train or operate each such system, (iii) any third-party providers or platforms on which each such system depends, and (iv) any known incidents, claims, or regulatory inquiries relating to any such system. To the Knowledge of the Company, each such system has been developed, deployed, and operated in compliance in all material respects with applicable Law, including Laws relating to data privacy, non-discrimination, and consumer protection.*"

   - **Purchase Agreement — Interim Operating Covenant**

     "*Between the date hereof and the Closing, except as set forth on Schedule [__] or as consented to in writing by Buyer, the Company shall not (a) deploy any new material artificial intelligence or machine learning system in the Business, (b) materially modify any existing AI system used in the Business, (c) enter into, amend, or terminate any material contract with a third-party AI vendor, or (d) submit any material proprietary data of the Business to a third-party AI system not previously disclosed to Buyer.*"

   - **Employment/Contractor Agreement — AI-Generated Work and Acceptable Use**

> "*Employee acknowledges and agrees that all work product created, developed, or generated by Employee in the course of employment, including any work product created with the assistance of artificial intelligence tools, shall be the sole and exclusive property of the Company. Employee shall not use any AI tool to perform services for the Company except as expressly authorized in writing by the Company's AI Acceptable Use Policy. Employee shall promptly disclose to the Company any use of AI tools in the creation of deliverables or work product.*"

- **Transition Services Agreement — AI Systems and Data Access**

  > "*During the TSA Term, Provider shall maintain and provide Recipient with continued access to the AI Systems and associated data sets identified on Schedule [__] (the "Transferred AI Systems") at service levels no less favorable than those in effect as of the Closing Date. Provider shall not modify, retrain, or degrade any Transferred AI System without Recipient's prior written consent. Upon expiration of the TSA Term, Provider shall deliver to Recipient all passwords, models, weights, training data, configuration files, and documentation necessary to operate the Transferred AI Systems independently, and shall permanently delete all copies of Recipient's data from Provider's systems.*"

AI risk is not a generic IT issue to be addressed with a single representation or a broad technology covenant. It is a cross-functional concern that touches confidentiality, intellectual property, regulatory compliance, employment, data governance, vendor management and post-closing operations. Buyers and sellers who rely on transaction documents drafted before AI became prevalent or treat it as an afterthought expose themselves to significant and potentially unquantifiable risk. Experienced M&A counsel should incorporate AI-specific provisions throughout the full transaction document suite, from the initial NDA through purchase agreements, employment agreements, transition services agreements and other ancillary documents. It is recommended that business management and owners consult with competent legal advisors on these matters.