

AI Deals Call For Tailored Approach To Address Hidden Risks

Published by *Law360*

June 18, 2026

Deal activity in the artificial intelligence space continues to advance at rapid pace and incredible values: On March 31, OpenAI secured \$122 billion in committed capital,¹ and on April 20, Anthropic PBC announced that Amazon.com Inc. is investing up to \$25 billion in the company to expand the existing partnership.²

In the U.K., with more than £6 billion (\$8.06 billion) of new investment in the AI sector announced in June by the government, including a £2 billion commitment from Advanced Micro Devices Inc. to accelerate AI innovation and research in the U.K. and a £1.7 billion commitment from Group NV to build out capacity in the U.K., the sector is likely to see a combination of new investment and mergers and acquisitions.³

Although the AI sector has seen some megadeals driven by large scale technology players, the AI deal universe also extends into the midmarket. And an AI deal is not just another technology deal. AI companies in sectors such as healthcare, legal services, financial services and logistics are increasingly attractive to both strategic buyers and financial sponsors.

In those deals, AI risk often sits alongside sector regulation, customer contracting, outsourcing, data protection and integration issues, making these sorts of transactions more than just a conventional software acquisition. Because of the nature of the asset class, an AI deal requires a different approach to diligence, risk allocation and deal execution. This article looks at each of those areas.

More Complex Diligence Questions

AI-related acquisitions are typically control transactions or strategic investments that have various commercial objectives, which often overlap. These include securing specialist talent, obtaining proprietary data, acquiring compute capacity, controlling distribution channels, and the race to own foundational model capabilities in the face of accelerating market consolidation.

Understanding these objectives is key because they determine what the buyer is paying for, and therefore drive where diligence and contractual protection should be concentrated.

Traditional diligence asks familiar questions: Who owns the assets, what liabilities exist, which contracts are material and what consents are required? AI transactions do not replace those questions; but they can make the answers harder to verify and more dependent on technical facts that legal teams must understand well enough to test. We consider some of these matters in detail.

Intellectual Property Ownership and Training Data

AI systems are trained on data and are not invented in the traditional sense, so the core intellectual property questions in an AI deal are seldom answered by a simple schedule of

registered rights. Model architecture, model weights, fine-tuning layers, prompts, evaluation data, deployment tools, open-source components, licenses for training data, and any third-party model or application programming interface dependencies will all be relevant.

The diligence exercise will therefore need to test whether the target had the rights it needed to build, train and commercialize the relevant IP assets. It will also need to test whether it can transfer them to the buyer. This can create genuine ambiguity around IP ownership, raising the question whether a model constitutes a patentable invention.

In most jurisdictions, current patent doctrine requires human inventorship, meaning that models themselves are not directly patentable, although the systems and methods used to build and deploy them often can be.

More practically, acquirers must examine how the target's training data was obtained, whether it was used under licenses that permit commercial exploitation, and whether open-source components were incorporated in ways that could infect proprietary model weights.

Copyright and database-rights risks should be addressed expressly where training or fine-tuning relies on scraped, licensed or customer-provided data. A buyer who inherits a model trained on contested data may also acquire significant latent liability.

The litigation landscape in relation to such issues is complex and evolving, so the share purchase agreement, or SPA, may need to elicit targeted disclosures and include specific warranties, special indemnities, escrow mechanics, purchase price adjustments or exclusions from warranty and indemnity insurance coverage.

Data Rights

In some AI transactions, the key asset may not be the model, but the dataset that made the model useful. Diligence should examine what data the target holds; how it was collected; whether it includes personal data; whether it was obtained under contractual restrictions; and whether it can be used for model training, fine-tuning, validation, product improvement and postclosing integration.

Change-of-control restrictions, purpose limitations, audit rights, data localization obligations and customer consents will all be relevant. In addition, the extent to which data is subject to the General Data Protection Regulation or the California Consumer Privacy Act, or equivalent regimes, will be key.

Finally, the existence or not of data processing agreements, licensing arrangements and consent frameworks may all contain restrictions on transfer or change of control that materially affect value and execution.

Talent Hires

Sometimes, AI acquisitions are done to secure people. Buyers may be acquiring a team of researchers, engineers and domain experts as much as a product or platform. Diligence should therefore cover invention assignment, confidentiality, consultancy arrangements, equity incentives, retention packages, restrictive covenants, immigration status, works council or employee consultation requirements, and the enforceability of noncompetes in the relevant jurisdictions.

Managing Regulatory Risk and Timing

AI M&A is subject to wide regulatory scrutiny across merger control, foreign investment, data protection, sector regulation and emerging AI-specific regimes. These issues affect deal structure, conditionality, long stop dates, interim operating covenants, and — potentially — break fees. So, they should be addressed as early as possible, ideally at the time of the letter of intent or term-sheet.

Competition Law

Antitrust authorities in the U.S., European Union and U.K. are focused not only on horizontal overlaps but also on control of key AI inputs: compute, data, models, talent and distribution. Practically speaking, this is likely to mean that an analysis of market share may be only the start of the regulatory analysis. Diligence may also need to consider exclusivity, cloud commitments, access to graphics processing units, preferred distribution rights and data access.

AI-Specific Regulation

The EU AI Act is now a central diligence item for EU-facing AI businesses, particularly where the target offers or uses systems in high-risk areas such as recruitment, credit, education, critical infrastructure or regulated services. Buyers should identify the state of technical documentation, conformity assessment obligations, governance arrangements and any compliance remediation required after closing.

Foreign Investment and National Security

AI is increasingly treated as strategically sensitive technology. The Committee on Foreign Investment in the United States, the U.K. National Security and Investment Act 2021 and similar regimes will likely be relevant where the target develops AI with defense, surveillance, cybersecurity, critical infrastructure, advanced semiconductor, dual-use or large-scale data applications. Filing requirements should be analyzed early and reflected in the SPA via conditions, cooperation covenants and the long-stop date.

SPA Provisions

AI company valuations can depend on technical claims that are difficult to verify: model accuracy, hallucination rates, latency, robustness, explainability, security, customer adoption and the ability to scale. To the extent possible, legal diligence should attempt to convert these claims into verifiable metrics that are covered in the SPA by warranties, indemnities or price-adjustment mechanics.

Earnouts and deferred consideration tied to technical milestones or product delivery are common in AI deals. They need very careful drafting, however, particularly where value depends on future technical milestones. The SPA will need to define the milestones and the measurement process carefully. To the extent possible, the agreement should also consider what happens if key personnel leave or changes in the product strategy after closing.

The SPA should address the identified AI risk profile rather than rely solely on generic technology warranties. Depending on the facts of the transaction, bespoke provisions may be needed for data rights, training data provenance, open-source compliance, third-party model dependencies, cybersecurity, model documentation, regulatory classification, material customer restrictions, employee invention assignments and postclosing access to technical personnel.

Assuming the deal is conditional upon regulatory clearances, interim covenants should be carefully considered. Between signing and closing, a target may retrain models, change datasets, update prompts, alter safety guardrails, enter into new cloud commitments, release new AI features or modify customer terms. If those actions could affect value, regulatory status or buyer integration, they should be addressed expressly in the conduct of business covenants and consent rights.

The availability of warranty and indemnity insurance should be considered early as part of the transaction structuring. Underwriters may ask targeted questions about training data provenance, open-source use, privacy compliance, cybersecurity, customer dependency, regulatory classification and model performance.

Coverage exclusions, and the pricing of AI-specific risks, including IP contamination from training data, data protection liability and model performance failures, are both evolving. The time to negotiate what is covered and what is excluded will need to be factored into the deal timetable.

Practical Checklist

It will be necessary to check whether the buyer is acquiring talent, data, customers, model capability, compute access, distribution or a regulated product. The answer should drive the diligence plan and the warranty package.

It will be important to identify proprietary model weights, fine-tuning layers, prompts, evaluation tools, inference infrastructure, open-source components, third-party application programming interfaces, cloud dependencies and deployment environments.

It is also necessary to ascertain who created the IP by confirming invention assignment, confidentiality, consultant ownership, retention arrangements, equity incentives, restrictive covenants and the enforceability of employee protections in each relevant jurisdiction.

Performance claims should be strictly determined by using definite milestones, governance, measurement rights, access to information, dispute resolution and personnel assumptions, and by specifying the consequences of product road map changes.

Broad warranties that the AI system is accurate or reliable should be narrowed and datasets, benchmarks, testing periods, error tolerances and exclusions for customer misuse or postclosing modifications should be defined. Finally, earnouts or deferred consideration metrics should be as specific as possible.

Check whether granular diligence will be needed on training, validation, customer and operational datasets, because data issues relating to provenance, licenses, personal data content, permitted purposes, transferability and exclusivity are likely to be significant issues.

A regulatory analysis should be done early on, to determine whether the transaction raises merger control, foreign investment, data protection, EU AI Act, sector regulation or sanctions issues, and build any required filings into the timetable and conditions.

The postclosing work required to remediate compliance gaps, migrate data, integrate cloud infrastructure, update customer terms, preserve key personnel and maintain model documentation should be identified prior to signing, to aid postclosing integration.

Customer contracts, data licenses, cloud agreements, research collaborations, university arrangements, government grants and consortium rules should be checked for change of control restrictions, consent requirements or restrictions on postclosing use.

Finally, before relying on insurance as the buyer's main remedy, it will be important to understand how underwriters view training data, open-source, privacy, cybersecurity, regulatory classification and performance claims.

Conclusion

AI M&A is not simply technology M&A with a new label. Lawyers advising on these sorts of deals need to be able to translate complex diligence findings into both the deal architecture and the deal documentation. Generic AI risk provisions are unlikely to be enough; certain key facts that move value will need to be understood.

These include data rights, technical dependencies, regulatory classification, talent retention, model performance and integration, and are likely to be what determines whether an AI acquisition delivers the asset the buyer thought it was buying.

[“AI Deals Call For Tailored Approach To Address Hidden Risks,”](#) *Law360*, June 16, 2026

**Subscription may be required for article access.*

1 <https://openai.com/index/accelerating-the-next-phase-ai/>

2 <https://www.anthropic.com/news/anthropic-amazon-compute>

3 <https://www.gov.uk/government/news/britain-powers-ahead-on-ai-with-billions-of-pounds-of-new-investment-and-thousands-of-jobs-secured-as-london-tech-week-wraps-up>

CONTACTS

For more information, contact your Katten attorney or any of the following attorneys.



Oliver Williams

+44 (0) 20 7770 5212

oliver.williams@katten.co.uk



Edward A. Tran

+44 (0) 20 7770 5254

edward.tran@katten.co.uk

Attorney advertising. Published as a source of information only. The material contained herein is not to be construed as legal advice or opinion.

©2026 Katten Muchin Rosenman LLP.

All rights reserved. Katten refers to Katten Muchin Rosenman LLP and the affiliated partnership as explained at katten.com/disclaimer.