THOMSON REUTERS
PRACTICAL LAW™

Acquiring an AI Company

by Annette E. Becker, Alex V. Imas, Jake Bernstein, Mark H. Wittow, Claude-Étienne Armingaud, Melanie Bruneau, Marion Baumann, Kenneth S. Knox, Julie F. Rizzo, Cameron Abbott, Thomas Nietsch, and Nicole H. Buckley K&L Gates LLP, with Practical Law Corporate & Securities

Status: **Maintained** | Jurisdiction: **United States**

This document is published by Practical Law and can be found at: **content.next.westlaw.com/w-040-1226**
Request a free trial and demonstration at: **tr.com/practicallaw-home**

A Practice Note highlighting issues to consider when counseling a prospective buyer of an artificial intelligence (AI) company. This Note discusses the primary due diligence issues relating to AI and machine learning (ML) and strategies to mitigate or allocate risks in the context of an M&A transaction. This Note is also helpful for AI company targets that seek to anticipate potential issues. In this Note, the term AI company refers to a company involved in the research, development, or monetization of a product or service that is primarily powered by an ML algorithm or model that creates functionality or utility through the use of AI.

Artificial intelligence (AI) is a field of computer science that aims to mimic human intelligence with machines, including the abilities to learn, reason, generalize, and infer meaning. These types of abilities are components of many AI tasks, including speech recognition, computer vision, translation, and generation of text and images. Recent advances in AI primarily involve generative AI which is a subset of AI that uses computer algorithms often employing large language models (LLMs) to generate outputs that resemble human-created content. But AI is not new. Its applications long pre-date Open AI's ChatGPT chatbot and DALL-E image generator. AI and machine learning (ML) have been in development since at least the 1950s, but have only recently advanced to the point of bringing AI into the public domain. For an overview of how AI and ML work, see Practice Note, Artificial Intelligence and Machine Learning: Overview. For more on LLMs, see Practice Note, IT Basics: Generative AI and Large Language Models: Overview.

Recent advances in AI, and in particular generative AI, have led to more merger and acquisition (M&A) activity in the AI sector. Buyers seeking to integrate AI into their products or services may find acquiring an AI company or its assets to be a more expedient way to own AI technology than building the technologies themselves. While there are many issues a buyer must evaluate when considering an AI M&A transaction, this Practice Note focuses on key issues for buyers, targets, and their counsel when planning an acquisition of or large investment in an AI company.

For this Note, an AI company is involved in the research, development, or monetization of a product or service that is primarily powered by an ML algorithm or model that creates functionality or utility through the use of AI. Companies that merely use AI or ML in some manner are not AI companies because almost every company does or will make use of AI in some form.

In general, the structure of an M&A transaction with an AI company target is similar to an M&A transaction with a target in another industry. However, M&A transactions with AI company targets introduce different regulatory and due diligence considerations, and so some modifications to the transaction documents and risk mitigating strategies are recommended.

To assist primarily buyers and their counsel, this Note identifies:

- Key due diligence issues for buyers to consider and address when evaluating an acquisition of an AI company.
- Mitigation strategies for limiting risks identified in due diligence.

This Note focuses on acquiring private AI companies, although the discussion regarding due diligence is also relevant to the acquisition of an AI company that is a public company. For convenience, this Note refers to the AI business being acquired as the "target," regardless of whether the transaction is structured as a purchase of

the equity of a private AI company or the purchase of the assets of an AI business.

While this Note primarily focuses on issues relating to acquiring a US company doing business solely in the US by a US buyer, for informational purposes it also includes a discussion of European data privacy and cybersecurity laws and European antitrust, merger control, and foreign direct investment considerations (see Box, EU and UK Considerations).

This Note does not address:

- All proposed or recently enacted law at the international, federal, state, and local levels that could impact AI businesses.
- All aspects of conducting an M&A transaction. This Note is intended to supplement other more general materials on M&A transactions. For non-industry specific private M&A resources, see Private Stock Acquisitions Toolkit and Asset Acquisitions Toolkit.
- Regulatory risk that may arise from the acquisition of AI companies operating in highly regulated industries such as healthcare, financial services, and certain manufacturing industries. For a discussion of AI in health care, see Practice Note, Artificial Intelligence for Health Care Providers: Overview.

For additional AI resources, see Artificial Intelligence Toolkit.

Due Diligence Issues

Buyers should conduct extensive due diligence to evaluate both the magnitude of the risks and the steps that they should take to allocate and mitigate acquired risks in the acquisition agreements.

The focus of this Note is on key diligence areas that are unique to an AI company. However, buyers should perform a comprehensive due diligence review covering all of the standard due diligence areas that are applicable to most companies generally, including organizational, capitalization, compliance with laws, financial, tax, litigation, and employee due diligence (see Private Mergers and Acquisitions Due Diligence Checklist).

Intellectual Property

Intellectual property (IP) legal due diligence of an AI company should follow the typical due diligence that is done for a technology-focused business such as a software, social media, or computer hardware company.

This due diligence typically involves a review and analysis of these areas:

- **Owned IP.** This includes identifying the target's proprietary IP, registered and unregistered, and proprietary software and other IT assets. It also includes understanding how the target's IP and IT assets may be encumbered or subject to restrictions or obligations that could adversely affect the buyer or the target's value.
- **IP-related agreements.** This includes a review of:
  - inbound license agreements relating to third-party IP and IT assets used by the target;
  - outbound licenses and customer agreements relating to the target's products and services (see Target's AI-Related Contracts); and
  - development, distribution, consulting, settlement, and other IP-related agreements.
- **IP disputes.** This includes a review of actual or potential IP-related disputes and assessing their potential impact on the target.

For more information on IP and IT issues to cover in legal due diligence, see IP Due Diligence Issues in M&A Transactions Checklist and Software, Cloud & Other IT Due Diligence in M&A Transactions Checklist.

AI company assets typically consist of software, computer systems, and materials used to develop AI technologies, primarily input data, training materials, and other materials of various types. Though these assets are similar in certain ways to other technology companies, there are differences. As a result, AI companies present unique IP due diligence issues that require attention in several key areas, including:

- Owning the AI (see Owning the Algorithm or Model).
- Permission to use the training materials (see AI Training Materials).
- Using AI-generated materials (see Authorship and Ownership).
- Protecting trade secrets (see Trade Secrets Policies and Practices).

Owning the Algorithm or Model

Every AI company owns or uses one or more AI algorithms or models. This business-critical asset is one of the main components of an AI-based system. However, many AI companies use another company's algorithms or models rather than owning them themselves. Buyers should carefully evaluate companies that do not own 100% of

their models or algorithms to assess the risks created by this lack of complete ownership.

For example, many companies purport to build products or services using OpenAI's application programming interface (API), but these companies do not own the underlying AI technology and their businesses are dependent on a third party's continued existence and development of its AI technology. Additionally, creating and maintaining certain kinds of AI products, such as LLMs, is so enormously expensive that smaller companies are usually unable to create them independently. These companies are called API-dependent AI companies to differentiate them from companies seeking to build their own wholly owned AI products and services. For more on APIs, see Practice Note, IT Basics: Application Programming Interfaces (APIs): Overview.

To help understand the algorithms or models to investigate, a buyer should request that the target identify:

- Existing AI products and services that are offered to its customers.
- AI products and services that are used internally.
- AI products and services that are in development.

A buyer should use this information to further identify and investigate potential IP and other related risks.

AI Training Materials

Another business-critical asset for an AI company is its training materials, also called training data sets, that are used to train the AI models. Training materials are an essential component of the ML process but are not simple to obtain.

Training materials may consist of data and databases, text or words, computer source code, images, video, or other types of materials. AI companies obtain and use training materials through a variety of methods, such as:

- Creating the data.
- Buying the data.
- Licensing the data.
- Tracking or modeling real-world observational data obtained from proprietary or publicly available sources.
- Web scraping or harvesting other publicly available data.

A key item for the buyer to confirm is that the target has the right to obtain and use the training materials for the development of AI products and services. The buyer should evaluate the risks arising from using training materials based on the source and nature of the training materials. Some primary risks relating to training materials are summarized in the following table:

Source of Training Materials	Potential Legal Issues
Owned	Privacy issues arising from the unauthorized use of personal data (see Data Privacy and Cybersecurity).
Licensed	• Contractual and copyright issues, such as: • Limitations on the scope of the license. For example, the licensed training materials: – cannot be used for AI development or training; or – must be used for non-commercial purposes only. • The licensor itself does not have sufficient rights to use or license the training materials. Licensors of training materials typically provide little if any in the way of representations or warranties regarding their own rights to gather and license the training materials
Obtained through web scraping or harvesting tools	Legal action by third parties, such as: • Under the Computer Fraud and Abuse Act (CFAA) for unauthorized access (however, these claims have mixed results, see Practice Note, Key Issues in Computer Fraud and Abuse Act (CFAA) Civil Litigation: Authorization Issues with Public Facing Websites). • Tort claims such as trespass to chattels or conversion. • Claims for unjust enrichment.

Several legal disputes concerning using various types of materials for training purposes raise both copyright and contract issues, including the following cases, each of which is at an early stage of the litigation process:

- *Doe 1 v. GitHub Inc.*, 2023 WL 3449131 (N.D. Cal. May 11, 2023). A California court rejected the defendants' argument that the plaintiff software developers had no standing to sue. The court instead determined that the software developers' property rights had been harmed and that there is a "realistic danger" that defendants' AI products will recreate the software developers' licensed code.
- *Andersen v. Stability AI Ltd.*, No. 3:23-CV-00201 (N.D. Cal.). Individual plaintiffs allege that defendants Stability AI, Midjourney, and DeviantArt use copyrighted images to train models for their AI image generation products without consent from or compensation to the plaintiff image rightsholders. The plaintiffs claim that the "new" images created by defendants' products are actually derivative composites. The output image is assembled from inputs of the artists' original works.
- *Getty Images (US), Inc. v. Stability AI, Inc.*, No. 1:23-CV-00135 (D. Del.). Plaintiff Getty alleges that defendant Stability copied millions of its photos without a license and used them to train its AI product, Stable Diffusion, to "generate more accurate depictions based on user prompts." Getty claims it licensed millions of digital assets to other leading technology innovators to train their AI models. Therefore, according to Getty, Stability infringes Getty's copyrights and engages in unfair competition.
- *UAB "Planner5D" v. Facebook, Inc.*, No. 19-cv-03132-WHO (N.D. Cal.). Plaintiff Planner 5D's website offers a home design tool which customers may use to design home interiors online. The tool uses object and scene files to train ML algorithms. According to Planner 5D, Princeton University used software to access hidden internet addresses where Planner 5D kept its object and scene files. The complaint further alleges that Princeton scraped Planner 5D's website to obtain its files and then shared them with defendant, Facebook. Planner 5D contends that, despite being publicly visible, underlying data, hidden internet addresses, and file locations are trade secrets unfairly obtained by Princeton.

For further developments in pending AI copyright and other cases, see Generative Artificial Intelligence: Federal Litigation Tracker.

Because training materials are typically reproduced, the reproduction, even if transitory, arguably is copyright infringement (17 U.S.C. § 106(1)). Fair use and de minimis defenses are available, but untested and unresolved.

For a discussion of fair use defenses, see Practice Note, Copyright Fair Use, 17 U.S.C. § 107, *Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith*, 598 U.S. 508 (2023), *Google LLC v. Oracle America, Inc.*, 141 S. Ct. 1183 (2021), and *The Authors Guild, Inc. v. Google, Inc.*, 804 F.3d 202 (2d Cir. 2015).

For a discussion of de minimis defenses, see Practice Note, Copyright Infringement Claims, Remedies, and Defenses: No Infringement, *Newton v. Diamond*, 388 F.3d 1189, 1195 (9th Cir. 2004), and *VMG Salsoul, LLC v. Ciccone*, 824 F.3d 871, 878-887 (9th Cir. 2016).

An argument could similarly be made that the AI model is copyright infringement because it is "based upon" the underlying training materials, and therefore a derivative work of the training materials (17 U.S.C. § 101 (defining derivative work as "a work based upon one or more preexisting works, such as a translation, musical arrangement, dramatization, fictionalization, motion picture version, sound recording, art reproduction, abridgment, condensation, or any other form in which a work may be recast, transformed, or adapted"); 17 U.S.C. § 106(2)). This argument seems unlikely to succeed, however, except in specific circumstances such as where the prompt directed the AI model to either one specific work, or a very small number of works.

Obtaining data from the public-facing internet via web scraping tools also raises issues beyond copyright, which are largely unresolved currently. Key sources of internet data may seek to limit access by imposing access controls and licensing requirements, and generally impose contractual restrictions that prohibit commercial uses. Those contractual restrictions are difficult to enforce broadly but may be successfully asserted against certain users or categories of users.

The laws governing the right to use publicly available materials for AI training purposes vary from country to country. Some jurisdictions have statutes that explicitly provide some rights to use publicly available materials for AI training purposes, including the UK, the EU, Singapore, South Africa, Australia, Thailand, and China.

Authorship and Ownership

To the extent the target is relying on AI-generated materials in its business, whether those materials consist of computer code, texts, images, videos, or music, the target may be unable to protect those materials from use and exploitation by others. The US Copyright Office has issued guidance generally rejecting claims of copyright authorship for works generated by AI tools

as distinct from human authors (see Copyright Office, Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence (Guidance), 88 Fed. Reg. 16,190 (Mar. 16, 2023)). For a summary of the Guidance, see Legal Update, Copyright Office Issues Registration Guidance on AI-Generated Content.

Copyright Office decisions rejecting copyright registrations for AI-authored works have generally been upheld in litigation (see *Thaler v. Perlmutter*, C.A. No. 22-1564 (BAH) (D.D.C. Mem. Op. Aug. 18, 2023); *Thaler v. Hirshfeld*, 558 F. Supp. 3d 238 (E.D. Va. 2021), aff'd 43 F.4th 1207 (Fed. Cir. 2022), cert. denied 143 S. Ct. 1783 (2023)). Similarly, the US Patent and Trademark Office (USPTO) has issued guidance indicating that patentable inventions must be created by human inventors, and the US Court of Appeals for the Federal Circuit has confirmed that AI-created inventions are not eligible for patent protection (see *Thaler v. Vidal*, 43 F.4th 1207 (Fed. Cir. 2022); Legal Update, Federal Circuit Confirms That Inventor Must Be Natural Person; the USPTO's AI-related patent resources).

Outside the US, some jurisdictions recognize copyright protection for AI-generated works, but many do not. The UK, Ireland, Hong Kong, India, New Zealand, and South Africa provide copyright protection for computer-generated works, while Australia, Germany, Brazil, Colombia, Mexico, and Spain require human creation.

Trade Secrets Policies and Practices

If the target provides AI products and services, the buyer should review the target's trade secret policies and practices. Other forms of intellectual property protection (patents, copyright, and trademarks) are not well suited to protecting the most valuable aspects of AI technologies, and targets most likely rely on trade secret rights to protect and maintain the value of the target's AI technologies. For general information about trade secrets, see Practice Note, Protection of Employers' Trade Secrets and Confidential Information.

Target's AI-Related Contracts

Buyers should review the target's agreements for AI products and services, regardless of whether the target is a party to the agreements as a customer or provider. In addition to standard contract diligence issues (see the bullet point on commercial contracts in Practice Note, Due Diligence for Private Mergers and Acquisitions: Categories of Materials and Common Issues), the buyer should consider how the target has allocated the various risks inherent in offering or using AI products and services between the contracting parties.

If the target is providing the AI product or service, the buyer should review:

- Whether the target's terms of service include or disclaim any representations and warranties regarding:
  - accuracy;
  - reliability;
  - non-infringement; or
  - use in sensitive areas such as health care or financial services.
- The indemnification obligations the target owes to its customers or the indemnification rights the target has against its customers.
- Whether the target limits its liability with damage caps and exclusions of certain damages types (such as consequential, indirect, or lost profits damages).
- The means of resolving disputes.
- Whether the target has ongoing obligations to provide support and updates.
- The use of open source software, in light of whether the AI technology is distributed to customers or provided as a hosted service. For a further discussion, see Practice Note, Open Source Software: Use and Compliance.

If the target is the user of third-party AI products and services, the buyer should review:

- The same issues listed above but viewed from the opposite perspective (for example, whether the third-party provider disclaims key representations and warranties).
- Whether the target uses third-party AI products and services in compliance with the contractual restrictions and in recognition of the limitations of that AI technology, including hallucinations and other errors.

Data Privacy and Cybersecurity

US federal and state level laws regulate collecting, using, processing, and disclosing personal information and address broad cybersecurity standards. Buyers should assess a target's data privacy and cybersecurity programs, giving special attention to the target's:

- Algorithms and models.
- Training data.

For guidance on addressing these issues in merger and acquisition transactions generally, including performing data privacy and cybersecurity program assessments, see Practice Note, Privacy and Data Security Due Diligence in M&A Transactions and Privacy and Data Security Due Diligence in M&A Transactions Checklist.

US Federal Data Privacy and Cybersecurity Laws and Regulations

Some federal laws regulating privacy and data security and collecting, using, processing, and disclosing personal information that may apply to AI products and services include:

- Section 5 of the Federal Trade Commission Act, which is a broad consumer protection law prohibiting unfair or deceptive trade practices that the Federal Trade Commisssion (FTC) has long applied to business practices that affect consumer privacy and data security.
- Sector-specific regimes such as the Gramm-Leach Bliley Act, which applies to financial institutions and the Health Insurance Portability and Accountability Act of 1996, which applies to most health care providers, health plans, and their service providers.
- Other laws that apply to certain activities or groups, such as the Children's Online Privacy Protection Act.

For more information on these and other federal privacy and data security laws that may apply to AI products and services, see Practice Note, US Privacy and Data Security Law: Overview: Federal Laws.

Federal regulators are also increasingly promulgating regulations regarding cybersecurity incidents, which may include personal data breaches and risk management. For example, the US Securities and Exchange Commission issued final rules in July 2023 requiring public companies to disclose:

- Cybersecurity incidents within four business days of determining materiality.
- Information regarding their cybersecurity risk management, strategy, and governance programs.

For more information, see Legal Update, SEC Adopts Cybersecurity Risk Management and Incident Disclosure Rules. AI companies often hold substantial amounts of data, so cybersecurity attacks on them can put a significant amount of data at risk. Buyers should therefore closely scrutinize how the target protects its data and assess the sophistication of the target's cybersecurity capabilities and processes.

US State Data Privacy and Cybersecurity Laws and Regulations

States have often taken a leading role in protecting their residents' personal information, for example, establishing data breach notification requirements and standards for reasonable data security practices. Similarities among these state laws exist, but there are also many differences that make nationwide compliance challenging.

Given the dependence of AI models' functionality on enormous amounts of data, often including personal information, the buyer must critically examine:

- The applicability of these state data privacy and cybersecurity laws to the AI company's business.
- The target's privacy and data security compliance programs, including any history of data breaches or related regulatory enforcement actions or private litigation.

More recently, the US has seen a rapid proliferation of consumer data privacy laws since California passed its first-in-the-nation California Consumer Privacy Act of 2018, as amended by the California Privacy Rights Act of 2020. Additional states have passed similar consumer data privacy laws, including Colorado, Connecticut, Delaware, Florida, Indiana, Iowa, Montana, Oregon, Tennessee, Texas, Utah, and Virginia, while others are considering them. These laws vary significantly in their scope and requirements, although they share some common themes nationwide and with the EU General Data Protection Regulation (Regulation (EU) 2016/679) (GDPR) (see Box, EU and UK GDPR Compliance).

These laws, and in some cases their implementing regulations, are important to AI companies that fall within their scope because they:

- Give residents of their states various rights to control uses of their personal information, including the right to opt out of certain types of data processing and to request deletion of their personal information under certain conditions.
- Impose various obligations on covered entities regarding how they communicate with individuals and process and protect personal information.
- Potentially limit AI companies from using their AI models to engage in automated decision-making, potentially impacting the value of those models to the buyer.

For more information on the state privacy and data security laws that may apply to AI companies, see State Data Privacy Laws Toolkit.

EU and UK Data Privacy and Data Protection Laws

M&A transactions may involve multiple jurisdictions, so non-US law should be considered as well. Most international jurisdictions have not created AI-specific laws or rules yet, but the EU and the UK have stringent personal data protection laws and some AI-specific laws to consider in an M&A context. For a further discussion and information on other global jurisdictions, see Box, EU and UK GDPR Compliance and Country Q&A Tool, Data protection.

Corporate Governance and Data Quality

Bias and Quality of Data

AI models are trained on existing data sets, which reflect societal biases, and can therefore continue to entrench these biases. As noted in the National Institute of Standards and Technology's (NIST) Artificial Intelligence Risk Management Framework (AI RMF 1.0) (AI RMF), AI "systems are inherently socio-technical in nature, meaning they are influenced by societal dynamics and human behavior." Because these models draw on often biased data, generate outputs that reflect these biases, and in turn incorporate those outputs into their training materials, these biases will tend to be magnified via adverse feedback loops if intentional remediation measures are not taken.

If an organization's data collection process systemically excludes or underrepresents members of certain groups based on race, ethnicity, gender identity, sexual orientation, or otherwise, training materials are likely to reflect and reinforce these biases. Similarly, if time or monetary resources are required to access data collection tools, training materials will exclude those who lack access to these resources. Understanding a target's data collection processes and reviewing them critically is a vital component of due diligence of an AI company.

When performing due diligence on an AI company, the buyer should seek to understand the impact that the target's organizational culture and employees have on the development of the system. An approach of "moving fast and breaking things" can create significant risk for AI companies in cases where haphazard practices could violate antidiscrimination and other laws. Due diligence processes should focus on how:

- Datasets are created or used by a target.
- The target's culture influences the generation of datasets.

Recently enacted laws have focused on bias at AI companies, and buyers should be aware of and ensure compliance with evolving state regulations. For example, New York City recently passed NYC Local Law 144 which took effect in July 2023 and requires companies using automated decision making technologies (ADMT) for employment decisions to conduct a bias audit on ADMTs within one year of using them and to notify candidates and employees of their use of such tools.

Internal Governance Policies

The buyer should also seek to understand the target's governance policies regarding data collection. Just as company boards of directors routinely undertake materiality assessments regarding environmental, social, and governance practices, boards can also mandate responsible use of AI tools.

Due diligence processes should focus on the steps the target's board has taken to ensure responsible use of AI and ethical gathering of data, including:

- Putting in place employee handbooks and training policies to train employees in using AI models responsibly.
- Identifying and mitigating bias.
- Putting in place an organizational culture that is attuned to discrimination and bias issues.

In addition to defining the organization's culture, it is important that the target's board put in place clear procedures regarding using AI responsibly. The more clearly defined acceptable use policies are, the more likely it is that a target will avoid relying on untrustworthy or inaccurate AI models. Failure to put in place appropriate measures can result in losses arising directly out of litigation, or indirectly out of loss of business from reputational damage.

At an organizational level, reviewing the target's risk management framework, such as whether it has implemented the NIST's AI RMF, can help to assess whether the organization aligns its values and goals in connection with its AI products. While NIST's AI RMF is voluntary, organizations that begin to implement the standards of this framework demonstrate a commitment to the responsible use and development of its generative AI models, and identifying whether targets use this framework provides useful signals to a buyer in assessing the target's risk profile.

For a discussion of using AI in the workplace and an example of AI workplace policies, see Practice Note, Artificial Intelligence (AI) in the Workplace (US) and Standard Document, Generative Artificial Intelligence (AI) Use in the Workplace Policy.

Antitrust and Merger Control Considerations

Antitrust considerations must also be part of the due diligence process. Buyers should evaluate the risk that an acquisition could be blocked through merger control procedures because it has an anticompetitive effect on the market or is cleared subject to the acceptance of commitments, restrictions (such as not relocating the target), or dispositions of assets. Competition authorities in recent years track with particular and increasing attention if and to what extent new forms of anticompetitive agreements and behavior are conceivable and taking place.

In nascent industries such as generative AI, antitrust risks are uncertain due to the lack of established dominant players. However, given the high profile of AI, the US FTC and the US Department of Justice (DOJ) may closely scrutinize M&A involving AI companies. The FTC, for example, has highlighted control of important AI inputs as a key part of their analysis, including over:

- Data.
- Talent.
- Computational resources.

M&A deals that consolidate these inputs or other critical or complementary applications in the hands of larger players may raise antitrust concerns. Acquisitions of nascent competitors may also raise substantive antitrust concerns. (See FTC: Generative AI Raises Competition Concerns (June 29, 2023).)

The agencies have been aggressive in the technology space generally. For example, they have recently sought to:

- Enjoin acquisitions by large market players of nascent companies on the grounds that these companies could feasibly develop competitive products internally. For example, in 2022, the FTC unsuccessfully sought to enjoin Meta's acquisition of Within Unlimited, Inc. on these grounds.
- Unwind acquisitions that they have since determined have an anticompetitive effect. For example, the DOJ is currently seeking to unwind Google's successful 2008 acquisition of DoubleClick.

Buyers should also review the target for potential antitrust liabilities (see Antitrust Due Diligence Checklist). The DOJ offers a safe harbor for any antitrust misconduct disclosed by the acquiring company within six months after the closing date of an acquisition (see Speech by Deputy Attorney General Lisa O. Monaco (October 4, 2023)).

M&A transactions may involve multiple jurisdictions, so non-US law should be considered as well. For a discussion of EU and the UK, see Box, EU and UK Considerations.

For a discussion of antitrust and merger control procedures in the US and an overview of the Hart-Scott-Rodino Act, see Practice Notes:

- Corporate Transactions and Merger Control: Overview
- Hart-Scott-Rodino Act: Overview.
- US Antitrust Laws: Overview.

Insurance Coverage

Buyers should perform due diligence on the target's insurance policies with a special focus on cyber insurance. Cyber insurance policies typically provide coverage for a variety of AI-related exposures, including privacy and copyright-related liabilities. Most cyber insurance policies provide coverage on a claims-made basis, which means the trigger for coverage is a claim made during the policy period. Many cyber policies also have change of control provisions that effectively terminate coverage when an acquisition occurs. Because cyber incidents may go undiscovered for months (or even years), buyers should consider purchasing tail coverage for post-closing claims that arise out of pre-closing events.

For general information on cyber insurance, see Practice Note, Cyber Insurance: Insuring for Data Breach Risk.

Mitigating Risks

Many of the risks discussed in this Note may be mitigated with diligent drafting of the definitive acquisition transaction documents, such as a stock or equity purchase agreement, merger agreement, or asset purchase agreement (generally referred to in this Note as acquisition agreements). There are several sections of the acquisition agreement where buyers can effectively allocate risks to the seller in ways that protect the buyer from significant liabilities. As used in this Note, the term "seller" refers to either the equity holders in the target (if the acquisition involves the acquisition of equity in private AI company or a merger of a private AI company) or the target (if the acquisition involves the sale of the AI assets of the target in an asset sale).

Representations and Warranties

Representations and warranties in an acquisition agreement help a buyer mitigate risk because they provide the buyer:

- With disclosure information to help identify risks.
- A basis for indemnification claims or representation and warranty insurance (RWI) claims if the seller breaches a representation (see Indemnification and Representation and Warranty Insurance).
- A basis for terminating the purchase agreement if the seller materially breaches a representation.

The typical representations and warranties used in M&A transactions offer a useful starting point for acquisitions of AI companies but require customization to address specific AI company risks. For an example of a non-industry specific purchase agreement with comprehensive representations and warranties that can be used as a starting point for customized AI representations, see Standard Document, Stock Purchase Agreement (Pro-Buyer Long Form). For examples of comprehensive representations addressing intellectual property and privacy and data security matters that can also customized for AI issues, see Intellectual Property in M&A Transactions Toolkit: Representations and Warranties.

To customize standard representations and warranties to an AI-related acquisition, the parties to the transaction should include an appropriately robust and accurate description of the AI technologies, products, and services that are the subject matter of the transaction. The target should make representations regarding:

- The AI technologies that it uses.
- What AI it offers as a product or service.
- What AI it plans to offer in the future.

Buyers should also seek to require the seller to bear the risks related to procuring and using training materials and claims related to the AI's outputs by including representations and warranties in the acquisition agreement confirming that:

- The target has identified or scheduled in the acquisition agreement all the sources of training data.
- The target has:
  - the rights to use each of those sources for training purposes; and
  - provided the buyer with the documentation establishing the target's right to use those training materials for training purposes.
- The quality of the sources of training materials is appropriate for their intended uses.
- The target has adequately designed and tested its AI products and services.

If the target provides AI products and services, the representations and warranties should confirm that:

- Those AI products and services:
  - have operated as intended;
  - have not been the subject of customer or third-party complaints or regulatory actions; and
  - are the subject of appropriately limiting contractual terms regarding appropriate use and limitation of liability.
- The target has not conveyed ownership or exclusive rights in any AI technology.

The representations and warranties should also address ethical and regulatory concerns by:

- Confirming that the target has appropriate guardrail policies and oversight in place.
- Requiring disclosure of complaints, investigations, proceedings, and litigation in this area.

If the target uses third-party AI products and services, the representations and warranties should:

- Require disclosure of all contracts for AI products and services.
- Confirm that the target:
  - is protected against liability for using the third-party AI products and services;
  - owns the output of its use of the third-party AI products and services; and
  - owns the IP rights in any third-party AI technology developed by a third party for the target, as with any customized technology paid for by the target.

For example AI representations by the seller, see the following Standard Clauses:

- Artificial Intelligence Representations: Asset Purchase.
- Artificial Intelligence Representations: Stock Purchase or Merger.

The buyer should also ensure the acquisition agreement contains robust representations and warranties regarding compliance with applicable laws, including data privacy laws, addressing both current and past compliance. For sample privacy and data security representations,

see Standard Clause, Privacy and Data Security Representations: Stock Purchase or Merger.

The seller, conversely, should try to limit the representations and minimize the seller's related liability exposure, by qualifying the representations with knowledge, materiality or MAE, and time period qualifiers. For an overview of ways to use these and other qualifiers to limit representations and warranties, see Practice Note, Stock Purchase Agreement Commentary: Limitations of Representations and Warranties.

Indemnification

Indemnification provisions provide a means of mitigating risks for a buyer in the purchase agreement by requiring the target to compensate the buyers for losses relating to specified issues arising prior to closing. Indemnification is commonly available to a buyer of a private company for losses resulting from inaccuracies or breaches of the seller's representations and warranties in the acquisition agreement.

Although, a seller's indemnification obligations in an acquisition agreement are often subject to survival periods, caps, and baskets that limit a seller's liability in the event of a breach of a representation, representations about fundamental matters are often carved out of these limitations.

Buyers should consider which representations and warranties may warrant a longer survival period or exclusions from the indemnification caps and baskets that are applicable to general representations and warranties. In particular, in the acquisition of an AI business, data privacy, intellectual property, and regulatory representations may merit special consideration as areas that could pose particular risks and therefore should be excluded from these limitations, including potentially via special indemnities. For a general discussion on indemnity limitations, see Practice Note, Indemnification Clauses in Private M&A Agreements: Limits on Indemnification Obligations.

In addition to indemnification for breaches of representations and warranties, a buyer may also negotiate for a special indemnity that obligates the seller to indemnify the buyer for losses arising out specified matters. These special indemnities may cover specific known risk areas including any losses arising out of pre-closing noncompliance with laws, pre-closing litigation, or any other high risk areas identified during the buyer's due diligence of the target.

Representation and Warranty Insurance

RWI in M&A transactions has become increasingly common as a means of allocating unknown risks with respect to breaches of representations and warranties from targets and minimizing the potential for post-closing disputes between buyers and sellers. For general information on RWI, see Practice Note, Representation and Warranty Insurance for M&A Transactions.

Third-party insurers are particularly aware of and sensitive to areas of risk in transactions involving AI companies. As the first generative AI company acquisitions have been completed, RWI insurers have noted that buyers seek to cover particular representations such as representations regarding sufficiency of assets, use of open-source software in company products, and data privacy similar to other technology M&A transactions. While, to date, underwriting practices and standards in transactions involving AI companies have not differed materially from transactions in the broader technology industry, insurers continue to be diligent in evaluating risks in AI companies, and their practices may evolve as more insured transactions involving AI companies are completed, and insurers continue to compile data on salient risks.

Closing Conditions

If there is a gap of time between signing and closing, each party to the acquisition agreement may require that the other party fulfill certain conditions before the transaction closes. Therefore, if the buyer identifies any significant issues and risks in due diligence, the buyer should consider requiring the target to make specified corrective actions as a closing condition. These may take the form of:

- Contractual amendments.
- Technology remediation.
- The withdrawal or cancellation of certain AI products or services.

If the seller does not satisfy the closing conditions, the buyer is not required to close and may terminate the transaction.

Post-Closing Remediation

Any significant issues and risks identified in due diligence that are not corrected before closing, should be addressed and corrected by the buyer after closing. These may take the form of contractual amendments, technology remediation, or the withdrawal or cancellation of certain

products or services. In the area of data privacy, for example, affirmative post-closing steps may be required to ensure ongoing compliance. In addition, if there are deficiencies in a target's policies and procedures, or no policies or procedures are in place at all, post-closing measures may be required to institute appropriate policies and procedures to ensure compliance with applicable laws and institute risk management practices that allow the business to limit liability and risk as it grows in an evolving legal environment.

For resources on drafting privacy and data security policies, see Privacy Compliance and Policies Toolkit.

EU and UK Considerations

In general, Europe is further along in implementing AI-specific legislation than other parts of the world. In the US, there are some guidelines and an AI-specific presidential executive order, but AI-specific legislation is largely limited to US state law. For a summary of the US presidential executive order on AI, see Legal Update, President Biden Issues Comprehensive Executive Order on Artificial Intelligence. Additionally, for more on developing US and key state AI regulatory developments, see Developments in US Artificial Intelligence Law and Regulation: 2023 Tracker.

Therefore, EU and UK considerations are relevant for AI M&A because:

- The target or a component of the target's business could be located in the EU or UK.
- The EU is ahead of the rest of the world in implementing AI-specific legislation. So, other jurisdictions may reference these EU rules when developing their own AI rules.

EU and UK GDPR Compliance

Companies doing business in the EU are subject to the EU GDPR, which governs personal data processing and took effect on May 25, 2018. The UK's post-Brexit retained version of the GDPR (UK GDPR) and Data Protection Act 2018 protect personal data in the UK. The UK GDPR currently remains substantially similar to the EU GDPR. Therefore, in this Note, the GDPR refers to the GDPR in the EU, the European Economic Area, and the UK.

The GDPR protects natural persons' personal data, including data AI companies may collect, such as:

- Information regarding an individual user's use of AI products.
- Information in the training materials used by the AI products.
- Users' inputs in response to prompts built into the AI products.

Compliance with the GDPR requires specific steps, including informing users of data processing operations and demonstrating a legal basis for personal data processing, which may include obtaining users' consent to make use of processed data for specified purposes. Failing to comply with the GDPR can result in substantial fines or injunctive relief, which may include requiring deletion of training materials or seizure of violating products. A buyer should ensure that the target has taken proper steps to comply with the GDPR, if applicable. For a further discussion of the GDPR, see GDPR Resources for US Practitioners Toolkit.

EU's AI Act

On June 14, 2023, the European Parliament passed an updated version of the EU Artificial Intelligence Act (AI Act), which governs providers, users, and producers of output of AI models that is used in the EU. Before becoming law, the European Parliament will negotiate with the Council of the EU and the EU member states. The AI Act is expected to categorize AI companies based on their evaluated risks. Similar to the GDPR, the AI Act creates an accountability framework for development and use of AI, including a requirement to document and assess AI risks and maintenance of certain transparency standards.

The current draft of the AI Act requires any entity controlling a high-risk AI system to:

- Publish a summary of the specific use and context in which the AI system is intended to operate (see Article 29, paragraph 6).
- Maintain a quality management system to ensure and document compliance with the AI Act (Article 17, paragraph 1, introductory part).

Buyers contemplating an acquisition of an AI company should ensure that targets are aware of the requirements of the AI Act and are taking appropriate steps to comply if it is applicable. For a further discussion of the AI Act, see Regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act): legislation tracker.

Antitrust Issues

In Europe, although there has not been a major case regarding AI merger control, there is increased regulatory scrutiny of digital markets and AI transactions by the European Commission (EC) and EU member states.

Even when merger control thresholds are not met at the EU or member states level, the EC may review the competitive effects of AI transactions when a member state makes an admissible referral (an Article 22 referral) (see Commission Guidance on the application of the referral). The newly interpreted referral mechanism described in the guidance targets "killer acquisitions" which involve targets active in highly innovative sectors, such as AI, whose current turnover is not representative of their competitive significance and which are therefore not captured by European or national filing thresholds.

Technology companies are under increased scrutiny when acquiring AI companies and may decide to notify the EC directly, which may consider the transaction to be a candidate for an Article 22 referral. The EC may also contact national authorities to invite them to make a referral and third parties may inform the EC of a transaction that could be the subject of a referral.

In the case of merger control procedures, buyers should evaluate the risk that an acquisition could be called in for review or even ultimately blocked (or be conditionally approved) because it has an anticompetitive effect on the market even in cases where the targets have no or limited turnover. Regarding new and innovative technologies, questions about access and foreclosure can play a decisive role in the assessment of potential anticompetitive effects (and in the interest of the EC to review a transaction).

Foreign Direct Investment

In the current political environment, foreign direct investment (FDI) considerations are even more crucial than merger control considerations. FDI regimes capture a potentially larger number of transactions due to different and broader triggering events.

In Europe, FDI screenings are made at the member states level and according to their respective national FDI laws. There is no standalone European FDI control regime. The EU Foreign Investments 2019/452 (FDI Regulation) does not include enforcement or veto rights for the EC but gives it the ability to issue opinions on certain investments. The FDI Regulation also seeks to harmonize the different screening mechanisms of member states while acknowledging member states' sole responsibility for their respective national security. The FDI Regulation explicitly addresses AI under Article 4(1)(b) as one of the "critical technologies" that are a factor to be considered by the EC and member states when assessing whether an FDI could affect security or public order. Accordingly, most member states have included AI as a sensitive sector in their respective FDI regimes.

In practice, this means that, if a target is active in a member state, it is necessary to determine

whether the acquisition must be reported to that member state as part of an FDI screening assessment. Twenty-one of the 27 EU member states now have a screening regime. FDI regimes usually require the presence of a local entity for them to be triggered but, in some cases, the presence of assets or sales can be sufficient. Buyers should be aware of the differing FDI regimes that apply in multiple jurisdictions.

For example, in Germany, the acquisition of a certain number of shares, a certain amount of control or of certain essential resources of a German entity by a non-EU member state would trigger an FDI notification requirement. In France, investments in or acquisitions of French entities active in certain sensitive sectors by foreign investors require prior authorization by the French Ministry of the Economy and Finance. In the UK, a blanket mandatory notification requirement is in place regarding investments in or acquisitions of UK companies engaged in certain AI-related activities by any investor, regardless of nationality.

This Note does not address FDI of foreign persons in US businesses. For general information on the power of the Committee on Foreign Investment in the United States (CFIUS) to review certain acquisitions of and investments in US businesses, see Practice Note, CFIUS Review of Acquisitions and Investments.