Digital Diligence: AI-Specific Diligence and Subject Matter Experts in Corporate Transactions

by Danny Tobey, Sean Fulton, and Coran Darling, DLA Piper, with Practical Law Intellectual Property & Technology

Status: **Law stated as of 02 Dec 2025** | Jurisdiction: **United States**

This document is published by Practical Law and can be found at: **content.next.westlaw.com/w-048-5534**
Request a free trial and demonstration at: **tr.com/practicallaw-home**

An Article providing a framework for conducting AI-specific due diligence in connection with a private company acquisition or investment. As AI becomes a key product or differentiator, it also acts as a risk multiplier, presenting unique challenges that traditional IP, privacy, and cybersecurity diligence may not uncover. This Article outlines preliminary scoping decisions based on transaction structure and the materiality of the target's AI development or use. It delves into the substantive areas of inquiry, including a target's proprietary AI development, use of third-party systems, training data governance, and the distinct issues raised by generative AI (GenAI). This Article also explores the rapidly evolving legal and regulatory landscape, from the EU AI Act to various US state laws, and emphasizes the importance of evaluating a target's internal AI governance program. By incorporating thorough AI diligence, acquirers and investors can better identify risks, build contractual protections, and confidently navigate deals involving AI technology.

In an era of expanding AI use and oversight, AI-specific due diligence is now a critical component of corporate transactions whenever the target company leverages AI in any meaningful way.

The first instinct in many transactions is to route everything that looks technical into familiar lanes, such as intellectual property (IP), privacy, cybersecurity, and commercial contracting, on the theory that these disciplines, taken together, capture most technology-related risk in modern businesses. While it is true that such traditional diligence may uncover some of these risks, it will not reliably surface the unique risks presented by the development, use, and deployment of AI, such as:

- Opaque model provenance.
- Tainted training data.
- Brittle evaluation and monitoring practices.
- Silent third-party model dependencies.
- Shadow deployments by business teams.
- Contract terms that quietly shift liability from vendors to the buyers.

On top of those somewhat static risks, AI also presents dynamic risks. AI systems learn from data, adapt to new contexts, and sometimes act differently when deployed than in testing environments. The accuracy of a system may degrade or drift without any code changes. The surrounding legal landscape is also moving. Risk derives not only from what the law prohibits today, but also what regulators and counterparties expect a responsible organization to do tomorrow. Such dynamic risks do not fit neatly inside older diligence patterns, and understanding these issues is critical to determining deal valuation, deal structure, risk allocation, integration plans, and post-close remediation budgets.

Because traditional due diligence approaches may not fully capture these AI-related concerns, it is becoming an imperative for buyers and their counsel to conduct specialized AI diligence to identify and mitigate these risks.

This Article outlines key considerations in establishing an AI diligence framework, including:

- Preliminary AI diligence scoping decisions (see Scoping AI Diligence: Preliminary Considerations).
- Substantive areas of inquiry into the target company's AI assets and uses (see Substantive AI Diligence).
- Relevant legal and regulatory risks (see Legal and Regulatory Landscape).
- The target's AI governance practices (see AI Governance and Organizational Framework).

Acquirers and investors who incorporate thorough AI diligence into their transaction playbook will be better positioned to close deals with confidence and integrate targets smoothly and safely (see Contractual Protections and Risk Allocation).

For an overview of considerations in M&A transactions involving AI companies, see Practice Note, Acquiring an AI Company.

For an overview of key legal issues relating to AI, see Practice Note, AI Key Legal Issues: Overview (US).

This Article focuses on acquiring or investing in private AI companies, although much of the discussion is also relevant to the acquisition of or investment in an AI company that is a public company. For an overview of due diligence in M&A transactions, see Practice Notes, Due Diligence for Private Mergers and Acquisitions and Due Diligence for Public Mergers and Acquisitions.

Scoping AI Diligence: Preliminary Considerations

Before the first virtual data room folder is opened, the buyer's team must determine the appropriate scope and focus of AI-specific diligence for the particular transaction. This initial scoping exercise ensures that resources are allocated efficiently to areas of highest risk or importance. By thoroughly addressing these preliminary considerations, the buyer's team will have a clear roadmap for the AI diligence process, including:

- What to focus on (see Key Factors Defining the Scope of AI Diligence).
- Where to find critical information (see Sources of Information).
- Whom to involve (see Subject Matter Expertise).

Key Factors Defining the Scope of AI Diligence

Transaction Structure (Investment Versus Acquisition)

The structure of the deal will influence the depth of diligence. For example:

- In a minority investment, the investor might perform a narrower review focused on specific concerns or representations.
- In a full acquisition of a target company, the buyer will typically require a comprehensive review of all AI assets, systems, and practices.

A minority investment deal may rely more on representations and warranties and ongoing covenants, while an acquisition (especially of an AI-centric company) demands exhaustive upfront diligence on technology and compliance.

Target's Commercialization of AI Versus Internal Use

Another crucial factor is whether the target sells or licenses AI solutions to customers (external commercialization) or merely uses AI internally to support its business.

A company that markets AI-driven products or services faces direct customer expectations and potential liabilities related to AI, requiring careful review of product performance, representations made to customers, and IP ownership of the AI inputs and outputs. In contrast, a company using AI only internally (for example, using AI tools to streamline its logistics) might present fewer outward-facing legal issues, though internal use can still raise data and compliance risks.

Extent and Materiality of AI Use

The team should assess how significantly the target relies on AI in its business. If AI technologies are central to the target's products, services, or operations, then AI-related diligence should be a major focus. Conversely, if the target uses AI only in ancillary ways (for example, a general retailer using a third-party AI tool for internal data analysis), the scope can be more limited. The more material AI is to the target's value proposition or revenue, the more extensive the diligence should be.

Role as AI Developer Versus AI Deployer

Counsel should consider whether the target is primarily:

- An AI developer, creating proprietary AI models or algorithms. Development raises issues related to ownership of algorithms and training data.
- An AI deployer, implementing third-party AI tools. Deployers will have issues related to vendor contracts, third-party dependencies, and compliance with licenses.

Understanding the target's role is also important for evaluating its compliance with requirements under laws and regulations that distinguish between AI developers (or providers) and deployers. For example, see Legal Update, Colorado Enacts Comprehensive AI Legislation Targeting Discrimination and Practice Note, EU AI Act.

Many companies are a mix of both, and diligence should tailor inquiries to each role.

Relevant Industries and Jurisdictions

The industry sector and geographic footprint of the buyer and the target can greatly affect AI risk exposure.

Industry-specific considerations are important. For example, an AI tool used in healthcare or financial services will attract scrutiny under healthcare privacy laws or financial regulations, respectively.

Jurisdictional considerations also come into play. Different countries and US states have emerging AI laws and regulations (see Legal and Regulatory Landscape). If the transaction is cross-border or the buyer or target operates internationally, then due diligence should account for compliance with each relevant jurisdiction's AI requirements (for example, the potential applicability of the EU AI Act to a US target selling AI products in Europe).

Buyers should also consider whether any export controls or trade restrictions apply to the target's AI technology, such as restrictions on exporting certain AI software or hardware to other countries.

Sources of Information

Effective AI diligence relies on gathering comprehensive information from several critical sources. It is essential to request and review documentation that the target company can provide about its AI systems, including:

- Model cards.
- Internal risk or impact assessments.
- Technical whitepapers.
- Testing and validation results.
- Audit results.
- Training data summaries.
- Policies or procedures governing AI development and use.

Such documents reveal how the company addresses issues like bias, privacy, and security. Notably, the absence of documentation signals immature AI governance and should itself be considered a diligence finding.

Key sources of information in public company due diligence reviews include publicly available information such as the target company's SEC filings.

Because of the highly technical nature of the review, direct interviews with the target's technical teams are oftentimes invaluable and can save resources when compared to traditional document review. These conversations often illuminate how AI tools are used in practice, uncover known limitations or incidents, and clarify future plans for AI.

Subject Matter Expertise

Engaging AI subject matter experts (SMEs), whether internal (such as the buyer's or target's Chief AI Officer or Chief Technology Officer or AI governance committee members) or external (such as outside counsel with appropriate expertise), can greatly enhance the diligence process. Their technical expertise and fluency can help bridge the linguistic gaps between legal and engineering teams. AI SMEs can:

- Evaluate algorithms and data from a technical perspective.
- Identify red flags such as opaque models or inadequate bias testing.
- Translate technical findings into business or legal risks.

Their expertise helps counsel ask the right follow-up questions and ensures that potential issues are properly understood and addressed.

Because AI-related risks often overlap with other areas of due diligence, coordination with specialist teams is essential. AI diligence typically involves:

- Privacy and cybersecurity specialists, if the target's AI systems process personal information or sensitive data, to ensure compliance and proper security controls.
- IP specialists, for review of inventions, patent filings, and third-party IP licenses and infringement risk.
- Employment law specialists, if the target uses AI in human resources functions.
- Consumer protection specialists, if AI outputs could impact consumers.

AI diligence is most effective as an interdisciplinary effort, drawing on all relevant domains of expertise within the diligence team.

Substantive AI Diligence

After determining the scope and assembling the right team, the buyer should conduct a deep dive into substantive areas of AI risk and compliance. Major categories of inquiry and concern that are unique to AI include:

- Proprietary development of AI technology (see Proprietary AI Development).
- Use of third-party AI technology (see Third-Party AI Systems).
- Deployment of AI technology in the target business (see AI Deployment).
- Training data (see Training Data and Data Governance).
- Generative AI (GenAI) use (see GenAI and Output Ownership).

Substantive AI diligence should be part of a comprehensive review covering all of the standard due diligence areas that are applicable to technology companies generally. For checklists covering typical areas of legal due diligence in M&A transactions involving technology companies, see:

- Software, Cloud & Other IT Due Diligence in M&A Transactions Checklist.
- IP Due Diligence Issues in M&A Transactions Checklist.
- Private Mergers and Acquisitions Due Diligence Checklist.
- Public Mergers and Acquisitions Due Diligence Checklist.
- Minority Investment Due Diligence Checklist.

Proprietary AI Development

If the target has developed its own AI models, algorithms, or software, the buyer should evaluate the ownership, quality, and risk profile of those assets.

Third-Party Dependencies

The buyer should identify any third-party components or dependencies within the target's AI. Many AI systems incorporate open-source libraries, pre-trained models (including large language models (LLMs) or other foundation models), or application programming interfaces (APIs) from external AI services. Due diligence should include a catalog of these dependencies and review of the license terms or usage rights. For open source software, the buyer should determine if the licenses are permissive or restrictive (for example, copyleft licenses could impose sharing requirements on derivative works). If the target's AI uses a third-party foundation model, the buyer should check what the usage terms are. For example, some foundation models have licenses restricting commercial use or requiring attribution.

The buyer should also verify that the target has the necessary rights to use any third-party datasets or pretrained models in its AI development. Unresolved third-party IP issues could pose significant risks post-acquisition, including potential infringement claims or unexpected costs to secure proper licenses.

Training and Retraining Processes

The buyer should examine how the target trains its AI models. Diligence should include requests for information on training datasets and methods used for training. The most important consideration often is whether the training datasets were obtained lawfully and with all necessary permissions, especially if personal information is involved. Other questions may include how the target ensures quality training data and the frequency of fine-tuning.

Consistent, well-managed training processes indicate a mature AI program, whereas ad hoc or one-off training might signal potential issues with model performance or data integrity.

Testing, Validation, and Readiness

A responsible AI development process includes rigorous testing and validation of models before they are deployed or used in production. During diligence, the buyer should request evidence of any model validation efforts. The absence of documented testing and validation is a red flag, as it may indicate the AI's efficacy or safety has not been verified.

Ongoing Monitoring and Maintenance

The buyer should determine how the target monitors the performance of its AI systems over time. Model drift (where a model's performance degrades as data patterns change) is a common issue. The buyer should ask if the target has monitoring in place to detect it. Lack of monitoring suggests that issues like accuracy degradation, bias creep, or security vulnerabilities in the AI might go unnoticed until they cause a problem.

Transparency and Interpretability

The buyer should assess the target's AI models' degree of transparency or explainability to mitigate the black box problem. In regulated industries, such as finance or healthcare, being able to explain an AI decision is crucial for compliance. If the target's model is highly complex (like a deep learning neural network), the buyer should ask what tools or methods the target uses to interpret model decisions (such as using explainability techniques or simpler surrogate models for explanation). A diligence report might note if a critical business decision relies on an AI whose decision-making process is opaque, as this could have regulatory or operational implications.

Documentation of the Development Process

Well-run AI projects usually maintain documentation on model design, assumptions, version history, and known limitations. The buyer should request any internal documentation or project notes on how the AI was developed. Documentation can reveal whether the target followed a structured process and considered ethical or legal issues during development (see Safety and Ethical Evaluations). It also helps the buyer later integrate or maintain the AI post-acquisition. A lack of documentation may mean knowledge is confined to certain key employees, posing a continuity risk if those individuals leave after the transaction.

Safety and Ethical Evaluations

The buyer should ascertain whether the target has performed safety assessments or ethical reviews if the AI technology could potentially cause harm or ethical issues. These may be appropriate, for example, where the target develops or uses AI that:

- Controls physical machinery.
- Makes impactful decisions about individuals.
- Can be repurposed maliciously.

Appropriate steps might include:

- Testing for unintended bias or disparate impact on protected groups.
- Evaluating worst-case failure modes.
- Convening an internal ethics committee to review the AI application.

Buyers are increasingly concerned with the ethical footprint of AI, not just legal compliance, because unethical AI use can lead to reputational harm and regulatory scrutiny.

Third-Party AI Systems

Many companies incorporate third-party AI systems or services into their operations rather than developing technology in-house. For example, a company may license a third-party's AI platform or engage a vendor to provide cloud-based AI services. If the target relies on third-party AI solutions, counsel conducting a due diligence review should scrutinize these arrangements and their implications.

Integration and Dependencies

The buyer should determine how the third-party AI tools are integrated with the target's own systems and data. If the target has built a product or service on top of a third-party AI platform, the buyer needs to understand the degree of dependency. High dependency means the target's business might be significantly affected by changes in the third-party's technology, pricing, or terms of service. The buyer should ask the target to document which parts of the target's operations would be impacted if the third-party AI became unavailable or if its performance changed.

Customizations

The buyer should check whether the target has made any custom modifications or add-ons to the third-party AI system. In some cases, a company might fine-tune a third-party model with its own data or build custom code around a third-party API.

These customizations can raise questions about ownership and maintenance, such as who owns the improvements or derivative works, and whether they will continue to function if the underlying third-party system updates or is no longer available. A diligence review of the vendor agreements should seek to determine what happens to customizations post-acquisition and whether the buyer will have rights to use or alter them going forward.

Review of Provider Contracts

A core part of this diligence is reviewing the contracts or terms of service with AI providers. Key clauses to scrutinize include:

- License scope and restrictions, to confirm whether the target has the rights it needs to use the AI for all intended purposes.
- Data usage rights, including whether the vendor can use the target's data for its own purposes (such as improving its AI technology).
- Risk allocation provisions, such as:
  - indemnities;
  - limitations of liability; and
  - warranties or disclaimers about the AI's performance.

Since AI products and services often involve nascent technology, vendors may offer services "as-is" without strong warranties. This is a risk area for the buyer to note.

Target's Diligence on Third-Party Providers

Counsel should inquire whether the target itself has performed any due diligence on the third-party AI system before adoption. If the target has copies of any vendor-provided compliance certifications, audit reports, or security assessments, those should be reviewed.

If the target did little to no vetting and just plugged a third-party AI into mission-critical processes, that may indicate an unchecked risk that the buyer will inherit.

Application of the Target's AI Policies to Third-Party Systems

If the target has its own internal AI governance policies or procedures, the buyer should confirm whether those policies extend to the use of third-party AI (see Application of Governance to Third-Party AI). Ensuring that third-party AI usage is consistent with the target's stated AI risk management practices is an aspect often overlooked, and any gaps might need remediation or representation in the deal.

AI Deployment

This category of issues focuses on how the target deploys and uses AI in practice, whether internally or in products and services, and what controls surround that deployment. Understanding deployment is critical because it affects the risk exposure and regulatory implications of AI.

Internal Use Versus External Use

The buyer should identify which AI systems are used internally (for example, for internal efficiency) and which are external or customer-facing. Internal uses of AI primarily raise concerns about internal governance, privacy, and employee impacts. External uses can create direct customer-related liabilities and reputational risks.

For external uses, the buyer should examine how the target's customers or other third parties interact with the AI and whether these parties are informed about AI involvement (transparency obligations), as well as any feedback or complaints the target has received.

High-Risk Use Cases

The buyer should determine if any of the AI deployments fall under categories considered high-risk from a legal or ethical standpoint. For each identified high-risk use case, diligence should probe deeper into the target's:

- Compliance with relevant regulations and industry guidelines.
- Implementation of adequate safeguards, which might include:
  - human review;
  - accuracy thresholds; or
  - bias mitigation.

Customer Contracts and AI-Specific Terms

If the target's AI is part of products or services offered to customers, the buyer should review the customer contracts for terms addressing the AI. Key clauses might include disclaimers of AI performance, limitations on use, and indemnification or liability clauses related to the AI's actions. The buyer may want to ensure that customer contracts are updated post-closing to better protect the company if needed.

Customizations and Configurations

Even for internally used AI or tools, the buyer should check if the target has customized them. For example, if the target has deployed an AI-based customer relationship management (CRM) plugin and configured it with certain datasets or rules, those configurations should be noted. Custom rules or AI model tweaks could themselves become important assets or potential liabilities if they cause AI to behave unexpectedly. The buyer should ask the target to confirm that knowledge of these configurations is documented and transferable.

Training Data and Data Governance

An AI system is only as good as the data on which it is trained. Thus, training data diligence is a vital component of AI-specific due diligence. The buyer should investigate the sources, legality, and quality of the data that the target's AI systems depend on.

Collection Methods and Data Sources

The buyer should ask the target to identify how the target obtained the data used to train its AI models. Common sources include:

- Data collected from the target's own users or systems.
- Purchased datasets from third parties.
- Web-scraped data.
- Publicly available datasets.

Knowing the provenance and contents of training data helps in assessing any legal exposure related to that data.

Data Preparation and Processing

The buyer should review what steps the target has taken to prepare the raw data for training. Good practices include data cleansing, labeling, anonymization or pseudonymization, and minimization. Proper data preparation affects not only model performance but also compliance. Inadequate anonymization or pseudonymization, for example, could mean the model is effectively trained on personal data subject to regulation.

Compliance with Licenses, Consents, and Permissions

The buyer should seek to confirm that for each dataset or data source, the target had the right to use it for AI development. This involves checking any license agreements or terms that cover the data. Non-compliance in this area can lead to major issues post-transaction, including lawsuits or the need to retrain models with properly licensed data, which can be very costly or impracticable.

IP Infringement Risk

Training data can pose hidden IP risks. For example, if the target scraped copyrighted content from the internet to train an AI model, the copyright holders of that content might have claims if such use is not protected by fair use or other exceptions. Diligence should attempt to uncover if any of the training data included content that might be protected by IP rights of others. The buyer might need to consider obtaining licenses, indemnities, or taking other steps to mitigate IP exposure.

Bias and Fairness Assessments

Closely related to data quality is the issue of potential bias in training datasets. The buyer should confirm whether the target has performed any bias audits or fairness assessments on its AI models (which inherently ties back to training data distribution). If the target has AI models making decisions about people (such as for employment or lending decisions), regulators and best practices increasingly call for bias testing. If bias issues are detected during diligence, the buyer will need to plan for mitigation (either pre-closing or as a post-closing integration task) to avoid liability or public relations fallout.

GenAI and Output Ownership

GenAI deserves special attention in due diligence because it introduces distinct issues regarding inputs, outputs, and contractual restrictions. For an overview of these issues, see Practice Note, Key Legal Issues in Using Generative AI: Overview (US).

If the target is using GenAI technologies (such as LLMs, image generators, or other AI that creates content), diligence should probe how these are used and managed.

Ownership and Rights in AI-Generated Outputs

The buyer should clarify who owns the outputs produced by GenAI and whether those outputs can be used freely by the target (and thus by the buyer after acquisition). Many jurisdictions currently treat AI-generated content as possibly not protected by copyright unless there is sufficient human creativity. If the target's business involves using GenAI to create valuable content (such as code, images, or reports), diligence should be conducted to determine whether the target can claim IP ownership of those outputs. If the value of the target hinges on AI-created assets, ensuring those assets are legally usable and ideally owned or exclusively controlled by the target is critical.

Output Consistency, Quality, and Accuracy

GenAI is known for occasionally producing incorrect results (hallucinations). The buyer should assess how the target manages the quality of AI outputs. The risk here is that unvetted AI outputs could mislead customers, contain inaccuracies that lead to poor decisions, or even include inappropriate or biased content. The buyer will want to know if any incidents have occurred and how the target responded. Diligence should include a review of the target's metrics on output accuracy or user feedback, if available.

Contractual or Policy Restrictions on Using Generative AI

The buyer should check whether the target is subject to any restrictions from third parties on using GenAI. A customer contract might prohibit the target from using GenAI to perform services, or a regulator or industry standard might discourage certain AI uses. For example, some financial regulators caution against using AI without certain controls, which could effectively bar unvetted GenAI in sensitive processes. The target's own internal policies might restrict employees from using generative AI for certain tasks or without approval.

The buyer should review any such constraints because they can directly affect how the target can operate. For instance, if key customers or a government contract disallows AI usage and the target was using it contrary to those terms, that may be a red flag. On the flip side, if the target has wisely restricted GenAI use in line with best practices, that's a positive sign of compliance and risk management.

Legal and Regulatory Landscape

AI is a rapidly evolving field not just technologically, but also in terms of law and regulation. An important part of AI diligence is identifying legal and regulatory risks that the target might face. These risks come from a patchwork of sources, from specific AI-focused regulations to general laws (such as privacy and consumer protection) that apply to AI uses. Buyers should be aware of the regulatory landscape during diligence so they can evaluate compliance and anticipate future obligations.

Developing AI Laws

The overall regulatory environment for AI is in flux. Multiple jurisdictions have enacted, or are proposing, AI-specific laws that affect how AI systems must be developed and used. A buyer should determine if the target is tracking relevant legislative developments and preparing for compliance. For example, the European Union's AI Act is a comprehensive regulation that imposes requirements (including transparency, risk assessments, and possibly registration) on certain uses of AI, with extraterritorial reach beyond Europe. For more information on the EU AI Act, see Practice Note, EU AI Act and EU AI Act Compliance for Non-EU Practitioners Checklist.

If the target's business might fall under a high-risk category (such as biometric identification systems or AI in recruitment), failing to plan for AI Act compliance could be a significant future liability.

State AI Laws in the US

Several US states have passed laws specifically addressing AI or automated decision-making. For instance, states like Colorado and Texas have comprehensive regulatory schemes addressing the use, development, and deployment of AI. Diligence should identify if the target or buyer operates in any jurisdictions with such laws and whether the AI systems in use are subject to them. If the state laws impose specific compliance obligations, such efforts should be reviewed. If the target has not accounted for these emerging requirements, the buyer will likely need to undertake compliance measures promptly.

For more information on recent AI legal and regulatory developments in the US, see Practice Note, Developments in US AI Law and Regulation: 2025 Tracker.

Consumer Protection Laws

The US Federal Trade Commission (FTC) and state attorneys general have signaled that they will use their authority over unfair or deceptive practices to police AI deployments that harm consumers. Diligence should consider whether any of the target's AI uses could be seen as harming consumers or whether the target has ever received regulatory inquiries or consumer complaints about its AI outputs. The buyer should also evaluate the risk of consumer protection actions and ensure the target's marketing of AI capabilities is truthful and not overhyped beyond what the AI can actually do reliably.

Product Liability Concerns

Traditional product liability law is now being applied to AI in some cases. Thus, when AI is embedded in products (software or devices) or services, the buyer should consider the risk of product liability claims or similar legal theories if the AI malfunctions. If the target's AI product could physically or financially harm someone if it errs, the buyer should ask how the target has insured or mitigated that risk.

For more information, see Practice Note, Artificial Intelligence and Tort Liability: The Evolving Landscape.

Sector-Specific Laws and Regulations

Many industries have their own regulations that can affect AI deployment. For example, in financial services, agencies like the Federal Reserve, the Office of the Comptroller of the Currency (OCC), and the Consumer Financial Protection Bureau (CFPB) have published guidance on the use of AI (especially around credit underwriting and fraud detection) that encourage transparency and fairness. In healthcare, the Health Insurance Portability and Accountability Act of 1996 (HIPAA) governs personal health information which might be used in AI, and the Food and Drug Administration (FDA) may regulate certain AI-driven medical tools or require validation studies. In insurance, state regulators scrutinize algorithmic underwriting.

If the target operates in a regulated industry, diligence should include a review of how the target's AI use fits within those specific regulatory frameworks. The buyer's counsel with sector expertise should be involved to evaluate compliance in applicable areas.

AI Governance and Organizational Framework

An increasingly important aspect of AI diligence is evaluating the target company's internal AI governance framework. Beyond looking at individual models or datasets, counsel should assess whether the target has an overarching system of policies, procedures, and practices to manage AI risk and compliance. A strong AI governance program at the target can significantly reduce the buyer's post-acquisition workload and risk, whereas the absence of one might mean the buyer has to implement governance from scratch. Depending on the transaction structure, the buyer may need to integrate the target's AI governance framework with the buyer's existing AI governance framework.

When assessing the target's AI governance, the buyer should consider whether the company has put in place:

- Formal policies (see Formal AI Policies).
- Employee training programs (see Training and Awareness Programs).
- Oversight mechanisms (see Oversight and Compliance Monitoring).
- Incident response processes (see Incident Response Process for AI Issues).
- Governance processes for third-party AI tools (see Application of Governance to Third-Party AI).

For a checklist that may be helpful when evaluating a target's AI governance, see AI Governance Checklist.

Formal AI Policies

The buyer should ascertain whether the target maintains written AI policies or guidelines. These could be standalone AI policies or part of existing policies (for example, information technology (IT), cybersecurity, privacy, or product development policies). A comprehensive AI policy might cover:

- Requirements for bias testing.
- Data handling in AI projects.
- Transparency standards.
- Approval processes for deploying new AI systems.
- Acceptable use of external AI tools by employees.

For a sample employee policy governing the use of GenAI tools in the workplace, see Standard Document, Generative AI Use in the Workplace Policy.

If the target has no formal policies, the buyer should assess how the target governs AI development and use. Sometimes smaller companies rely on informal norms, which might not be sufficient as they scale. The buyer should gauge policy maturity, as it indicates the level of seriousness with which the target approaches AI risks internally and may indicate the likelihood that a risk materializes.

Training and Awareness Programs

The buyer should confirm whether the target has provided training to its employees (both technical and non-technical staff) about AI policies, ethical AI practices, and regulatory compliance. Effective AI governance often starts with ensuring that personnel are aware of the company's AI standards and the legal implications of AI use. Diligence might include requests for:

- Evidence of training sessions.
- Internal AI best practices documentation.
- Designated AI ethics officers or champions within the organization.

Oversight and Compliance Monitoring

The buyer should ask the target to identify who, if anyone, is responsible for overseeing the target's AI compliance. Robust governance might involve an AI committee or working group, possibly drawing members from legal, compliance, IT, cybersecurity, privacy, and business units, that reviews new AI initiatives and monitors existing ones. Alternatively, some companies designate a specific executive (such as a Chief AI Ethics Officer) or incorporate it into the Chief Data Officer's role to keep track. If the target is large enough, internal audit or compliance departments might have started to include AI in their audits.

The presence of a robust oversight mechanism will reduce the risk of unnoticed issues, while its absence means the buyer might have to institute one and could discover latent issues that were never monitored.

Incident Response Process for AI Issues

A good governance framework will include a plan for what to do if something goes wrong with an AI system, such as discovering bias, a significant error, or a security breach involving an AI model. The buyer should inquire:

- Whether the target has faced any AI-related incidents or complaints.
- How the company responded to the incident or complaint.
- If no incidents have occurred yet, who would handle it.

It's important for a buyer to know that if an AI-related problem emerges, the company has the agility and process to address it promptly to minimize legal fallout or harm.

Application of Governance to Third-Party AI

As part of governance, the buyer should confirm whether the target's framework extends to employee and vendor procurement and use of external AI tools. Many companies now require some form of vetting or approval before a new AI tool is adopted, especially if it will handle sensitive data or critical functions. If the target has a vendor management or procurement policy, the buyer should confirm whether it identifies AI-specific procedures, such as requiring vendors to answer a questionnaire about their AI's risk measures or requiring legal sign-off if an AI will process personal data. This aspect of governance is increasingly crucial because an otherwise compliant company can introduce risk simply by plugging in a risky third-party tool without review.

Contractual Protections and Risk Allocation

After identifying the various issues and risks through AI diligence, the buyer and its counsel should consider how to address these findings in the transaction documents. The acquisition or investment agreement can be crafted to protect the buyer against AI-related risks (for example, by shifting risks to the seller or ensuring the target takes certain corrective actions before closing).

Examples of contractual protections and mechanisms that parties use to allocate and mitigate AI-related risks include:

- Representations and warranties, and related disclosure of AI issues in disclosure schedules.
- Indemnification provisions, including:
  - special indemnities for specific AI matters; and
  - carve-outs from general survival periods, baskets, and caps for AI-focused representations.
- Representation and warranty insurance.
- Pre- and post-closing covenants to address AI issues.
- Closing conditions and related termination rights.

For further discussion of these AI-related risk mitigation strategies, see Practice Note, Acquiring an AI Company: Mitigating Risks.

For sample AI representations, see Standard Clauses:

- AI Representations: Asset Purchase.
- AI Representations: Stock Purchase or Merger.

For a discussion of AI representations in recent publicly filed M&A agreements, including links to sample language, see Practice Note, What's Market: AI Representations in M&A Agreements.