EPIC Insurance Brokers & Consultants

PROFESSIONAL, EXECUTIVE AND CYBER SOLUTIONS

Beyond the Sandbox:
Governance Focused AI Contracting

Since the launch of ChatGPT in November 2022, the corporate world has moved from initial curiosity to experimentation and now to rapid, widespread implementation of various Artificial Intelligence (AI) and Generative AI (GenAI) tools. In 2026, AI and GenAI tools have transitioned from a conceptual "transformational force" to a critical operational and competitive necessity. While previous years saw AI and GenAI largely confined to innovation labs and experimental sandboxes, 2026 represents a definitive pivot point with Agentic AI making its bold entrance as well.

As organizations shift from experimental pilots to scaled, autonomous systems, a rigorous review of AI vendor contracts must become a key risk management function. AI tools and services present unique legal, ethical, and operational risks that need to be understood and proactively addressed.

Below are a few key contract clauses to look out for and consider:

1. **Subprocessor Clause:** First and foremost, it is important to remember that many AI vendors build their products on top of underlying third-party Artificial Intelligence models or services like OpenAI, Anthropic, AWS, etc. An AI Subprocessor Clause regulates how your vendor uses, accesses, and processes your organization's data through these third-party AI tools. Some key elements to look for in an AI Subprocessor Clause include:

   (a) an explicit list of all authorized subprocessors, including their names, locations, and the specific services they provide together with a requirement they provide you advance notice of any new subprocessors;

   (b) a vendor requirement to impose the same data protection, security, and confidentiality obligations on its subprocessors as those imposed on the them in your contract;

   (c) a primary vendor obligation can bear liability for all obligations it delegates to its subprocessors, including any acts or omissions by the subprocessor that violate data privacy laws;

   (d) prohibiting the use of your data for training or improving the subprocessor's models – this may or may not be achievable but in either instance it should be considered as a risk; and,

   (e) a clause allowing you the right to request a copy of the agreement between the vendor and the subprocessor to verify the required data security standards are met. This is certainly not an exhaustive list, but these elements should provide a good jumping off point for evaluating the clause.

2. **Data Usage and Model Training Clause:** Many AI vendor contracts allow vendors to use customer data for retraining models or for competitive intelligence purposes. A data usage and model training clause in an AI vendor contract is a provision that either explicitly permits or prohibits a vendor from using a customer's data—including inputs, prompts, outputs, and metadata—to train, improve, or fine-tune their AI models, algorithms, or related services. Consider whether your vendor contract **permits** or **prohibits** the AI vendor from using your organization data or outputs generated from that data to train, retrain or improve their AI model. If the clause permits the vendor to use your organization's data, consider whether that is in line with your existing obligations to your organization's clients, customers, vendors, business partners or employees. If that permission conflicts with your existing obligations, you may need to redraft this clause to **prevent** the use of data (inputs, outputs, prompts, etc.) for model training, retraining, fine tuning, or product improvement.

3. **Model Drift and Maintenance Clause:** Unlike static software, AI models can become less accurate or develop biases as real-world data changes with the passage of time. A model drift clause in an AI vendor contract is an important provision that requires the vendor to monitor, detect, and remediate the decline in performance or accuracy of an AI model over time. Some AI vendor contracts will treat model drift as a performance degradation event requiring a mandated remediation plan. A "retrain-in-30-days" clause can help ensure the vendor maintains accuracy of the model over time. Review your contract to determine how your AI vendor addresses model drift and how that position could impact risk for your firm or organization.

4. **Intellectual Property and Data Ownership Clause:** Data ownership and intellectual property (IP) clauses in AI vendor contracts are critical for determining who owns input data, model outputs, and any improvements or "learnings" generated during the contract term. Unfortunately, IP ownership is complex and cannot simply be divided into a straightforward "you own it" or "they own it" arrangement. When reviewing IP and data ownership clauses, prioritize clearly defining ownership of "Background IP" (pre-existing) vs. "Foreground IP" (newly created), ensuring strict confidentiality, and auditing indemnification scope. Key risks to consider include inadvertent IP transfer, loss of data control, and infringement lawsuits. Consider whether your current insurance portfolio captures and would respond to IP related lawsuits and resulting financial loss.

5. **Transparency Clauses:** Transparency clauses typically require AI vendors to disclose how their AI systems are developed, trained, and operated, including disclosing model limitations, data sources, and intended use cases. These clauses help firms and organizations comply with regulations (like the EU AI Act) by providing the necessary documentation to audit data, audit model performance, and ensure accountability. They enable organizations to identify and mitigate risks surrounding algorithmic bias, data privacy, and ethical compliance, ensuring a more defensible framework for their AI deployment. Omitting a transparency clause in an AI vendor contract creates a "black box" scenario where your organization lacks visibility into how the AI functions, what data it uses, and how it reaches decisions. This lack of oversight introduces several critical legal, operational, and financial risks.

6. **Regulatory Compliance Clause:** Regulatory Compliance clauses in an AI vendor contract are essential for mitigating the unique legal, ethical, and operational risks associated with AI technology. Currently, AI technology is evolving at a pace that often exceeds the legislatures' ability to address the risk.

Global regulations like the EU AI Act, domestic state-level laws like the Colorado AI Act and local laws such as the NYC Local Law 144, are emerging rapidly and with no level of consistent requirements. Regulatory compliance clauses serve as a critical defense against potential regulatory fines/penalties, intellectual property disputes, and data privacy violations. Some elements to consider in these clauses include:

(a) a statement that the vendor's AI system, training data, and outputs comply with current regulations and a commitment to meet future regulatory change;

(b) the right to audit the AI vendor to verify compliance;

(c) a provision requiring the AI vendor notify customers of regulatory inquiries, investigations, or substantial breaches; and

(d) an indemnity obligation wherein the AI vendor accepts liability if their AI tool causes your organization to violate legal obligations.

7. **"As is" or No Warranty Clauses:** AI vendor contracts frequently use "as-is" or "no warranty" clauses to avoid liability for hallucinations, inaccuracies, and biased outputs, transferring the risks of using Generative AI to the purchasing organization. One potential counter to this type of clause is to negotiate specific Service Level Agreements (SLAs) connected to measurable output quality and performance metrics. Since AI is probabilistic, a strict "zero error" SLA is likely unrealistic. Instead, focus on measurable accuracy metrics, traceability, and human-verified workflows that connect and support your use cases.

8. **Liability and Indemnification Clauses:** As Generative AI (GenAI) and agentic AI tools move from proof-of-concept to production, liability and indemnification clauses are rapidly evolving from traditional Software-as-a-Service (SaaS) models toward specialized, "shared responsibility" frameworks. AI risks often sit between the provider and user, leading to a split liability structure. Due to the high-stakes nature of AI—including hallucinations, data privacy breaches, algorithmic discrimination, and intellectual property (IP) infringement—contracts now commonly demand higher liability caps ("super-caps") or uncapped indemnity for specific breaches. For example, negotiating separate, higher liability limits specifically for data privacy and regulatory fines (like EU AI Act violations) is becoming the new baseline for enterprise contracts. Agentic AI tools introduce even more complexity. Consider elements such as requiring detailed audit rights for agent actions and ensuring the indemnity extends to third-party claims arising from the agent's autonomous performance. Bottom line, these clauses are a critical aspect of risk management and transfer. They are evolving rapidly and need to be evaluated carefully with experienced advisors.

9. **Insurance Clauses:** Insurance is often a backstop to indemnification. Because AI can create large-scale liability, insurance requirements should be considered carefully and aligned with the indemnification clauses, to the extent possible. The insurance market is beginning to introduce AI-related exclusions and limitations in commercial insurance products. As a result, most enterprise buyers of AI products or services are starting to require AI vendors carry more specialized insurance policies such as Technology Errors & Omissions (Tech E&O) and Cyber Insurance; often requiring a combined "Tech E&O and Cyber" policy. The question of what limits of insurance to require is challenging given the rapidly evolving nature of technology and the regulatory environment. Consider the potential "worst-case" downstream financial loss rather than just the contract value and consult with an experienced insurance broker or advisor.

10. **Termination and Transition Clauses:** Termination and Transition clauses serve as your contract's "exit strategy," ensuring business continuity and preventing vendor lock-in when parting ways with an AI vendor. One thing to look for in these clauses is a requirement for the return of your data in a machine-readable format or a “Certificate of Destruction” (CoD) (a formal, auditable document that provides verified proof that your sensitive data has been irreversibly destroyed or permanently erased). Another area of focus is on a transition "bridge period” which is typically around 90 days. The bridge period is essential to secure technical support for a seamless migration to a new vendor. Finally, look for wording that revokes the vendor's access to live systems post-termination and clarify whether these transition services are included in your base fee or billed at additional hourly rates.

AI vendor contracting is distinctly different from traditional Software-as-a-Service (SaaS) contracting and requires a significantly higher focus on governance and risk management. While traditional SaaS agreements focus on stability and uptime, AI vendor contracts must address the probabilistic, opaque, and rapidly evolving nature of AI technology. The above provisions are just a few of the critical clauses that you should consider.

Transitioning toward governance-focused AI contracting is a strategic necessity today that can transform your AI vendor contracts from passive documents into active risk-management tools. A robust, forward-looking contract—reinforced by comprehensive cyber and professional liability insurance—is essential to codify accountability, transfer residual risk, and safeguard against failures in complex AI-driven ecosystems.

Ready to Become EPIC?

For more information about how EPIC can help your business, visit us at **epicbrokers.com**.