

SEPTEMBER 30, 2026 | PUBLICATION

Client Alert: Artificial Intelligence Broke the Software Contract: The New Battle Over Indemnities, Warranties, and Liability Caps

For decades, software contracts followed a relatively predictable pattern.

Customers pushed for broad intellectual property (IP) indemnities, meaningful warranties, and liability caps large enough to provide a practical remedy if things went wrong. Vendors responded by narrowing obligations, limiting damages, and carving out as much risk as possible. The parties negotiated, found a middle ground, and closed the deal.

Artificial intelligence (AI) has disrupted that balance.

Not because traditional contract principles no longer apply; they do. The challenge is that many of those principles were developed around deterministic software. If a software product malfunctioned, the parties generally had a clear understanding of what happened, why it happened, and who was responsible. AI systems do not always provide the same level of certainty.

As a result, some of the most contentious negotiations in technology transactions today are no longer focused on functionality or service levels. They are focused on a much more fundamental question:

Who bears the risk when an AI system produces an undesirable result?

The answer increasingly depends on how the parties negotiate indemnities, warranties, and liability limitations.

Software Was Predictable. AI Is Probabilistic.

Traditional software is designed to perform specific functions according to defined rules. If a payroll system calculates taxes incorrectly or a financial application corrupts data, the error can generally be traced to a coding issue, configuration problem, or operational failure.

INDUSTRY SECTOR

Technology

SERVICE LINE

Data Privacy
Technology

RELATED PROFESSIONALS

Lloyd J. Wilson

MEDIA CONTACT

Wendy M. Byrne
wbyrne@shumaker.com

Generative AI operates differently. The same prompt may produce different responses. Outputs may be inaccurate, incomplete, biased, misleading, or entirely fabricated. Most vendors now openly acknowledge that AI-generated content should be reviewed by humans before being relied upon.

That disclaimer is understandable from a technical perspective, but it becomes much more complicated from a contracting perspective.

Many customers purchase AI-enabled products precisely because they are expected to assist with decision-making, content generation, analysis, forecasting, or operational efficiency. If the customer is paying for those capabilities, it is reasonable to ask what happens when they fail.

That question is now moving to the center of contract negotiations.

The Quiet Retreat from Traditional Warranties

Historically, software vendors were often willing to provide warranties stating that the product would perform substantially in accordance with documentation.

Increasingly, vendors are attempting to carve AI functionality out of those commitments.

Common formulations include language stating that:

- AI-generated outputs may contain errors.
- Outputs are provided on an “as is” basis.
- The vendor makes no representation regarding accuracy, completeness, reliability, or fitness for a particular purpose.
- Customers are responsible for validating all AI-generated content.

From the vendor’s perspective, these provisions reflect technical reality. Large language models are not designed to guarantee perfect accuracy.

From the customer’s perspective, however, these disclaimers can significantly reduce the practical value of the product.

A customer purchasing AI-assisted contract review software, threat-detection tools, financial forecasting capabilities, or compliance-monitoring solutions may reasonably ask why it should bear the full risk of inaccurate outputs generated by the very feature it is paying to use.

The result is a growing disconnect between product marketing and contractual commitments. The sales presentation emphasizes intelligence, automation, and enhanced decision-making, while the contract simultaneously disclaims responsibility when those capabilities fail. General counsel are increasingly recognizing this tension.

The New Fight Over Intellectual Property Indemnities

Few provisions have generated more debate in AI transactions than IP indemnification.

Traditional software indemnities were generally straightforward. If the software infringed a third party’s IP rights, the vendor would defend the claim and assume responsibility for resulting damages, subject to negotiated limitations.

AI complicates that framework. Questions that rarely arose in traditional software transactions now appear routinely:

- What if an AI-generated output allegedly infringes a third party's copyright?
- What if a model was trained using data that becomes the subject of future litigation?
- What if the customer incorporates AI-generated content into commercial products?
- Does the vendor indemnify the output, the model, both, or neither?

Many vendors now attempt to limit indemnity obligations by excluding claims arising from customer prompts, customer-provided data, customer modifications, or use of generated outputs.

Customers often respond that these exclusions are swallowing the indemnity entirely. After all, if prompts and outputs are excluded, the risk allocation may provide little practical protection for the activity most likely to generate disputes.

The underlying issue is not merely legal. It is economic. Someone must absorb the risk associated with unresolved questions surrounding AI-generated content. Vendors and customers are now negotiating intensely over who that party should be.

Liability Caps No Longer Reflect the Potential Exposure

Liability provisions are also under pressure. Most software agreements were built around assumptions regarding foreseeable damages and relatively predictable failure scenarios. AI introduces the possibility of losses that are harder to quantify.

Consider a hypothetical AI-enabled compliance platform that incorrectly classifies regulatory obligations. Or, an AI security tool that fails to identify a threat. Or, an AI system that generates inaccurate information that becomes embedded in customer-facing materials. The consequences may extend well beyond licensing fees.

Customers increasingly argue that traditional liability caps based on annual subscription costs may not reflect the actual risk profile of AI-enabled products. Vendors, unsurprisingly, take the opposite position. Many are attempting to preserve conventional liability structures while simultaneously expanding warranty disclaimers and narrowing indemnification obligations.

The practical effect is that customers may bear a greater share of AI-related risk than they historically assumed in technology transactions.

The Emerging Middle Ground

Despite the rhetoric surrounding AI, most negotiations ultimately land somewhere between the extremes. Rather than insisting on unlimited liability or blanket disclaimers, sophisticated parties are beginning to focus on targeted risk allocation.

Examples include:

- Separate indemnity provisions addressing AI-generated content.
- Enhanced transparency regarding training data practices.
- Representations concerning compliance with documented AI governance programs.
- Obligations to conduct testing, monitoring, and model oversight.
- Specialized liability caps for particular categories of AI-related claims.

- Requirements for human review in high-risk use cases.

These approaches are not perfect. They do, however, reflect an important shift. The discussion is moving away from whether AI creates risk and toward how that risk should be allocated. That is ultimately the purpose of a commercial contract.

What General Counsel Should Be Asking Now

As organizations continue deploying AI-enabled products, several questions deserve attention during negotiations:

1. **What warranties actually apply to AI functionality?**
2. **Does the IP indemnity cover AI-generated outputs, and if not, why not?**
3. **What exclusions materially reduce the value of the indemnity?**
4. **How does the vendor address model training, testing, and governance?**
5. **Do liability caps align with the risks the technology is expected to create?**
6. **Are we assuming responsibility for risks the vendor is in a better position to manage?**

These questions are becoming increasingly important because the traditional software contract was built for software that behaved predictably. Many AI systems do not.

Conclusion

AI has not eliminated traditional contracting principles. Indemnities, warranties, and liability limitations remain the primary tools for allocating risk between parties. What has changed is the nature of the risk itself.

The assumption that software produces consistent and predictable results no longer applies in every circumstance. AI systems may generate value precisely because they operate differently from traditional software. That same characteristic, however, creates uncertainty regarding accountability when things go wrong.

As a result, technology transactions are entering a new phase. The central issue is no longer whether a company will adopt AI. Most already have. The real negotiation is determining who bears the consequences when the model is wrong.

If you have questions or would like more information about indemnities, warranties, and liability caps in AI contracts, please contact Lloyd Wilson.