Shumaker

JUNE 12, 2026 | PUBLICATION

When Your SaaS Vendor Becomes an AI Company Overnight

Most companies did not sign up for an artificial intelligence (AI) relationship with their software vendors. They signed up for a platform that processed data according to defined rules, produced predictable outputs, and operated within the boundaries of a negotiated agreement. For a growing number of organizations, that is no longer the case, and in many instances, the change has occurred without any corresponding change to the governing contract.

INDUSTRY SECTOR

Technology

SERVICE LINE

Data Privacy
Technology

RELATED PROFESSIONALS

Lloyd J. Wilson

MEDIA CONTACT

Wendy M. Byrne
wbyrne@shumaker.com

Over the past 12 to 18 months, Software as a Service (SaaS) providers across virtually every category have integrated AI into their platforms. Some did so with fanfare, marketing new AI-powered features as a reason to renew. Others did so quietly, embedding machine learning models, generative tools, or algorithmic decisioning into existing workflows through routine product updates, often without renegotiating a single contract term.

For in-house counsel, this presents a problem that is less about technology and more about governance. The vendor you evaluated, conducted due diligence on, and contracted with may no longer resemble the vendor currently operating inside your environment. The agreement governing that relationship may not account for what the platform is now doing with your data, your users' inputs, or the outputs it generates on your behalf.

The Quiet Transformation

The pattern has become consistent enough to warrant a plain description. A company procures a SaaS platform for a specific function: customer relationship management, human resources, document management, contract lifecycle management, or financial reporting. The agreement includes standard provisions addressing data use, confidentiality, security obligations, and intellectual property (IP) ownership. At the time of execution, the platform operates as traditional software. Inputs are processed according to defined logic, and outputs are deterministic.

Then, through a product update or a revised terms of service, the vendor introduces AI capabilities. A recommendation engine begins surfacing suggestions. A generative drafting tool appears in the interface. Predictive analytics start influencing workflow prioritization. In some cases, the vendor begins using customer data, including inputs, usage patterns, and content, to train or refine its AI models.

These changes often arrive without a contract amendment, without updated data processing terms, and without formal notice that the nature of the platform has materially changed. The vendor's terms of service may contain a clause reserving the right to modify the platform or to use aggregated or anonymized data to "improve services." Language that once covered routine service improvement may now authorize model training. What appeared unremarkable in purely deterministic software now carries materially different implications when the platform is learning from your data.

The Federal Trade Commission (FTC) has taken notice of this dynamic. In guidance issued in February 2024, the Commission warned that companies that quietly change their terms of service to expand data practices, particularly to support AI training, risk engaging in unfair or deceptive conduct under Section 5 of the FTC Act.[1] The agency emphasized that a business that collects data under one set of privacy commitments cannot unilaterally adopt more permissive practices through a retroactive amendment without first providing prominent notice and obtaining affirmative consent. Where companies have previously represented that data would be used for limited purposes, expanding those purposes to include model training without adequate disclosure may trigger enforcement.

Why This Is a Governance Problem, Not Just a Procurement Problem

The consequence is straightforward: most legal teams are already carrying AI exposure embedded in existing SaaS agreements that were never designed to address it. The instinct for many legal teams is to treat this as a contracting issue, one that can be resolved by tightening the next round of vendor agreements. Better data use restrictions, explicit AI training carve-outs, and stronger audit rights are all important, and I will address them below. Focusing only on future agreements, however, misses the more immediate exposure: the SaaS contracts already in your portfolio.

In a previous article, I discussed how 72 percent of S&P 500 companies now disclose AI-related risks in their U.S. Securities and Exchange Commission (SEC) filings, up from just 12 percent in 2023, with disclosures increasingly focused on reputational harm, cybersecurity vulnerabilities, and regulatory scrutiny.[2] Those disclosures are not merely abstract. They reflect a growing recognition that AI risk is embedded in vendor relationships, not just internal deployments.

Among the S&P 500 companies disclosing cybersecurity risks tied to AI, a significant number have flagged third-party and vendor risk as a core concern, emphasizing that even strong internal safeguards cannot offset exposure if critical vendors are compromised. This observation is directly relevant to the SaaS transformation problem. When a vendor integrates AI into a platform that handles sensitive business data, the customer's risk profile changes regardless of whether the change was requested or formally approved.

The governance gap becomes particularly acute when vendor-side AI models process personal data, generate content that could carry IP implications, or make recommendations that influence business decisions. In each of these scenarios, the customer organization may bear regulatory, contractual, or reputational consequences for outcomes shaped by a system it did not evaluate, did not approve, and may not fully understand.

The risk is not theoretical. A human resources platform that uses employee data to refine its models may create privacy exposure that was never disclosed to employees. A generative drafting feature embedded in a

contract management tool may introduce third-party content into outputs, raising latent IP issues. A Customer Relationship Management (CRM) recommendation engine that prioritizes certain customers over others could create bias or discrimination concerns if it influences downstream decision-making. In each case, the customer organization bears the consequences of outcomes shaped by a system it did not meaningfully evaluate.

What In-House Counsel Should Do Now

The practical response requires both remediation of existing agreements and preparation for future ones. Neither is optional if the goal is defensible governance.

1. **Identify where AI already exists.** Work with IT and procurement to flag vendors that have introduced AI features since contract execution. This includes reviewing product release notes, updated terms of service, and privacy policy revisions. The objective is to build a clear picture of where AI is now operating inside your environment and whether the current contract adequately addresses it.
2. **Map data flows to AI functionality.** For each vendor that has introduced AI, determine what data the AI component accesses, whether customer data is used for model training or improvement, and whether outputs generated by the AI are treated differently from traditional software outputs under the existing agreement. Pay particular attention to broadly drafted "service improvement" or "aggregated data" clauses that may now encompass AI training activities that were not contemplated at the time of execution.
3. **Assess contract gaps.** Review existing agreements against the specific risks that AI introduces. Key questions include whether the agreement addresses ownership of AI-generated outputs, whether it restricts the vendor's ability to use customer data for model training, whether it includes representations about AI accuracy or bias testing, and whether indemnification provisions cover AI-related claims such as IP infringement arising from model outputs.
4. **Engage vendors directly.** Where contracts are silent or ambiguous on AI, initiate conversations with vendors about their AI practices. Ask what data is used for model training, whether customers may opt out, what safeguards exist to prevent data leakage across the model, and how the vendor handles AI-related incidents. Vendors that have adopted AI responsibly will be able to answer these questions with clarity and specificity. Vague or incomplete responses should be treated as a risk signal, not a neutral outcome.
5. **Update your vendor governance framework.** In my earlier article on AI governance, I recommended that legal teams create a centralized registry of all AI systems, including those embedded in third-party tools. This recommendation becomes more urgent as vendors integrate AI without formal notice. The registry should capture not only purpose-built AI tools but also AI features embedded in platforms originally procured for non-AI functions.

Strengthening Future Agreements

For new SaaS agreements or renewals, counsel should consider incorporating several provisions that address the AI transformation risk directly.

First, require affirmative notice before the vendor introduces AI features that process customer data. A general reservation of rights to modify the platform is insufficient when the modification fundamentally changes how data is used. The agreement should define what constitutes a material change and require advance written notice with an opportunity to evaluate and, if necessary, terminate.

Second, include explicit restrictions on the use of customer data for AI model training. The default should be that customer data is not used for training purposes absent express written consent. Where training is

permitted, the agreement should specify the scope, the anonymization or aggregation standards applied, and the customer's right to revoke consent.

Third, address IP ownership for AI-generated outputs with specificity. Traditional SaaS agreements typically address ownership of customer data and vendor IP in the platform itself. AI introduces a third category, namely outputs generated by the vendor's AI using the customer's data, that existing IP provisions may not adequately cover.

Fourth, require representations and warranties addressing AI-specific risks, including bias testing, accuracy limitations, and compliance with applicable AI governance frameworks such as the NIST AI Risk Management Framework or ISO/IEC 42001. These representations should be supported by audit rights that allow the customer to verify compliance.

Finally, ensure that indemnification provisions expressly cover AI-related claims. As AI models generate content, recommendations, or decisions, the risk of third-party IP infringement or regulatory violations shifts in ways that traditional indemnification language may not capture.

The Broader Point

The transformation of SaaS vendors into AI companies is not inherently problematic. AI-enhanced platforms can deliver meaningful improvements in efficiency, analytical insight, and operational capability. Those benefits, however, do not eliminate the obligation to govern the relationship with the same rigor applied to any other enterprise risk.

The challenge for legal teams is that this transformation often happens incrementally and without formal acknowledgment. A feature update here, a revised privacy policy there, and over the course of months, the platform bears little resemblance to the one originally evaluated and approved. When the vendor changes, the governance should change with it.

For general counsel, the question is not whether your SaaS vendors are using AI. Many already are. The more pressing question is whether your agreements, your oversight mechanisms, and your governance frameworks have kept pace. In the new AI era, the issue is no longer whether vendors use your data but whether your contracts still control how they do.

For more information on the impact of AI on SaaS vendor contracts, please contact Lloyd Wilson or another member of Shumaker's Technology, Data Privacy, Cybersecurity & AI Service Line.

[1] Staff in the Office of Technology and The Division of Privacy and Identity Protection, *AI (and Other) Companies: Quietly Changing Your Terms of Service Could Be Unfair or Deceptive*, Fed. Trade Comm'n (Feb. 13, 2024), https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/02/ai-other-companies-quietly-changing-your-terms-service-could-be-unfair-or-deceptive.

[2] Lloyd J. Wilson, *From Disclosure to Defense: A Strategic AI Governance Blueprint* (2025), https://www.shumaker.com/insight/from-disclosure-to-defense-a-strategic-ai-governance-blueprint/.