Tucker Ellis LLP

The AI Vendor Contract Is Becoming More Important Than the Privacy Policy

AUGUST 2026

Overview

Most companies today are not building artificial intelligence – they are buying it. As organizations across industries increasingly procure AI capabilities from third-party vendors rather than developing them internally, the locus of privacy and legal risk has shifted to the vendor relationship—and the contract that governs it.

For years, the privacy policy was the centerpiece of data governance. It remains relevant, but it no longer tells the full story. Today, the AI vendor agreement is the primary instrument through which companies control (or fail to control) the legal, regulatory, and reputational risks of AI deployment. These are not routine technology agreements; they are critical governance documents and should be negotiated accordingly.

This Client Alert identifies key provisions that in-house counsel and business teams should prioritize when evaluating and negotiating AI vendor agreements.

Key Contractual Provisions

Use of Customer Data for Model Training. This is perhaps the single most important provision. The contract should expressly prohibit the vendor from using customer data to train, retrain, or improve its models, whether using the customer's own model or the vendor's general-purpose models, without the customer's explicit, informed consent. Without this restriction, proprietary data could be absorbed into models that serve competitors or commingled with other customers' data, with no practical way to retrieve, segregate, or remediate it afterward.

Data Retention and Deletion. The agreement should specify clear retention periods for all categories of customer data, including prompts, outputs, inference logs, and metadata, and, where appropriate, impose affirmative deletion obligations upon contract termination or expiration. Vague commitments to delete data "in due course" or "within a reasonable period" are insufficient.

Subprocessors. AI vendors routinely engage subprocessors for cloud hosting, model inference, and data annotation. The contract should require advance disclosure of, and

approval rights over, all subprocessors and mandate that equivalent data protection obligations flow down to each one. Companies should also negotiate the right to object to new subprocessor appointments.

Security Incidents. Notification timelines are critical in the event of a breach or security incident. The agreement should require the vendor to notify the customer within a defined and compressed window (ideally 24 to 48 hours of discovery), preserve relevant forensic evidence, cooperate fully with the investigation and remediation, and assist with required regulatory notifications.

Confidentiality. Standard confidentiality clauses often fail to account for the unique nature of AI interactions. The agreement should explicitly protect prompts, outputs, and any business logic or proprietary information embedded in customer queries from unauthorized disclosure or reuse.

Audit Rights. Companies should secure meaningful audit rights, including the ability to conduct or commission independent third-party assessments of the vendor's data handling practices, security posture, and compliance with contractual and regulatory requirements.

Data Location. For companies subject to the General Data Protection Regulation (GDPR), sector-specific regulations, or data localization requirements, the contract should specify where data is stored, processed, and accessed. It should also restrict cross-border transfers absent appropriate legal safeguards, such as standard contractual clauses or adequacy determinations.

AI-Generated Decisions. Where AI outputs inform or drive consequential decisions, such as those involving hiring, lending, underwriting, or benefits administration, the contract should require the vendor to provide meaningful model explainability, support human-in-the-loop oversight, accept defined accountability for errors or bias in automated outputs, and support compliance with applicable regulatory frameworks, including the EU AI Act.

Indemnification. Thoughtful risk allocation is essential. The vendor should indemnify the customer for losses arising from data breaches, intellectual property infringement (including claims that the vendor's training data or outputs infringe third-party rights), and regulatory violations attributable to the vendor's platform or practices. These critical categories of risk should be carved out of broad liability caps and exclusions.

Ownership of Prompts and Outputs. The agreement should clearly vest ownership of all customer prompts and AI-generated outputs in the customer. Any license the vendor retains, even for limited operational purposes such as service delivery or troubleshooting, should be narrowly scoped, time-limited, and explicitly stated.

Regulatory Cooperation. As regulatory scrutiny of AI systems intensifies at both the federal and state level, vendors should contractually commit to cooperating with the customer's responses to regulatory inquiries, audits, and investigations. This includes preserving relevant documents and complying with reasonable litigation-hold obligations.

Implications

Privacy policies remain a necessary element of any data governance program, but they are unilateral statements, not negotiated commitments. The vendor contract is where enforceable rights and binding obligations are established. For in-house counsel and procurement teams, rigorous negotiation of AI vendor agreements is essential to responsible AI governance. Key takeaways:

- Treat AI vendor agreements as strategic governance documents, not administrative formalities. Companies that do so will be better positioned to manage the legal, regulatory, and operational risks of AI adoption.
- Prioritize the provisions identified above—particularly restrictions on model training, data retention, and indemnification—during contract negotiations.
- Revisit existing AI vendor agreements to ensure they reflect current best practices and emerging regulatory requirements.

Tucker Ellis can assist your organization with the review and negotiation of AI procurement contracts and help you manage the legal, regulatory, and operational risks that accompany AI adoption.

Additional Information

For more information, please contact:

- Paul J. Malie | 216.696.3466 | paul.malie@tuckerellis.com
- Anthony R. Petruzzi | 216.696.5478 | anthony.petruzzi@tuckerellis.com