

[The Licensing Journal, Negotiating AI Provisions in Commercial and Technology Contracts: Where the Market Is Heading, \(May 1, 2026\)](#)

The Licensing Journal

[Click to open document in a browser](#)

Negotiating AI Provisions in Commercial and Technology Contracts: Where the Market Is Heading

Doneld Shelkey

Doneld Shelkey is the global leader of Morgan Lewis's technology transactions, outsourcing, and commercial contracts practice. He represents clients across a wide array of complex and strategic transactions involving IP, technology, and other commercial matters. Don regularly negotiates deals that focus on AI, data rights, assets, software as a service, cloud computing, data acquisition and services, enterprise resource planning, blockchain, software licensing and development, social media, sponsorships, and strategic joint ventures and business collaborations.

Two years ago, many technology agreements addressed artificial intelligence (AI), if at all, through a generic disclaimer or a brief acknowledgment that AI features might be included in the offering. Today, that approach is inadequate. The integration of AI into commercial products, outsourcing arrangements, and enterprise software agreements has forced a rethinking of longstanding contract frameworks.

The SaaS Framework Is Under Pressure

Technology agreements for AI products have largely been structured as software-as-a-service (SaaS) arrangements: the customer pays for platform access, the vendor maintains the model, liability is capped, and the customer bears responsibility for how it uses the tool. That framework made reasonable sense when AI was a passive capability with a human making all consequential decisions.

However, the emergence of agentic AI and deeply embedded AI functionality is creating tension as the realities of “sufficient human intervention” become more apparent. We are seeing sophisticated buyers push for service-oriented terms: defined service descriptions, performance-based warranties, governance and audit rights, and outcome-tied liability structures. Vendors are resisting some of these provisions, but the negotiating dynamic is slowly shifting, at least in deals where AI performs critical enterprise tasks.

Key Negotiating Points

Scope and Service Definitions

Where AI performs tasks on a customer's behalf, the service description needs to articulate what the AI is actually doing—the workflows it executes, the decisions it makes, the systems it touches, and the guardrails that apply. Vague scope definitions create gaps that both parties will regret when something goes wrong.

Performance Standards and Warranties

Standard SaaS agreements for AI products frequently include broad disclaimers discarding any warranty as to accuracy or fitness for purpose. Where AI is embedded in core business processes and outcomes matter, buyers are increasingly unwilling to accept a warranty-free posture. We are seeing more deals include commitments around performance benchmarks, accuracy thresholds (which, although time-intensive to create, are increasingly important), and obligations to remediate when outputs fall short.

Intellectual Property Ownership and Output Rights

Intellectual property (IP) allocation in AI agreements has become one of the most actively negotiated areas. Key issues to consider include who owns AI-generated outputs, what rights the customer retains in training data it provides, and whether the vendor can use customer data to improve its model. We recommend addressing output ownership, training data use limitations, and data return and deletion rights explicitly—tailored to the actual use case rather than carried over from a generic SaaS template.

Regulatory Compliance Allocation

The regulatory landscape for AI varies widely—from the EU AI Act to US state-level requirements, such as those in Colorado, to federal guidance in financial services and healthcare. Agreements should address which party bears responsibility for compliance, how obligations are monitored, and how liability is allocated for regulatory penalties arising from the AI's use. Neither party benefits from ambiguity.

Transition and Termination

End-of-term planning should address continued access during wind-down, extraction of fine-tuned model components or proprietary configurations, return or deletion of customer data, and transition assistance. These provisions are often negotiated with less attention than they deserve—until a customer seeks to change vendors.

Avoid the Generic AI Disclaimer Trap

The better approach for both buy-side and sell-side is to integrate AI risk into the substance of the contract, including in the IP provisions, the liability framework, the data governance terms, the compliance allocation, and the service-level commitments. A well-drafted AI agreement reflects the specific use case and the actual risk profile of the deal. While it may require more work upfront, it produces an agreement that functions when it matters.