

The Licensing Journal, Contractual IP Protection in the Age of AI: Structuring Deployment and Development Agreements, (Oct. 1, 2026)

The Licensing Journal

Contractual IP Protection in the Age of AI: Structuring Deployment and Development Agreements

Matt Kohel

Matt Kohel represents clients in commercial litigation, intellectual property matters, and data privacy issues. His experience with intellectual property matters includes claims for the misappropriation of trade secrets, trademark infringement, false advertising, and the sale of counterfeit goods. In addition, Matt counsels companies that develop and use artificial intelligence (AI) on AI governance and strategy, intellectual property, data privacy, and regulatory and compliance issues. He writes regularly and is a frequent speaker and panelist on AI, especially as related to intellectual property, data privacy, and corporate policies on the use and management of AI. Relying on his experience as a litigator, Matt drafts and negotiates agreements and licenses for businesses, with a focus on protecting intellectual property and proprietary technologies.

I. Introduction: The End of the Traditional Software Boilerplate

For decades, enterprise technology procurement has relied on a predictable set of contractual assumptions. Standard Software-as-a-Service (SaaS) agreements, Master Services Agreements (MSAs), and IT licensing frameworks were built for deterministic systems, which are environments where static software code processes defined inputs to produce predictable, repeatable outputs. In these traditional software paradigms, intellectual property (IP) risk was largely bounded, data flows were predictable, and boilerplate contract terms governing data ownership and use provided adequate protection.

The deployment of enterprise artificial intelligence (AI) systems, and in particular, generative AI (GenAI) models, upends this traditional framework. Modern AI architectures do not operate as static tools. Instead, they engage in dynamic model interactions, exhibit non-deterministic behaviors, and generate fluid data loops. There is inherent unpredictability in the manner in which GenAI systems create output. When an organization deploys a Large Language Model (LLM), integrates Retrieval-Augmented Generation (RAG) architectures, or implements autonomous agentic workflows, this unpredictability blurs the boundary lines of proprietary data, derivative works, and output ownership.

This technical evolution collides directly with a stark legal reality. In particular, under current federal doctrine, statutory intellectual property regimes do not protect AI-generated expressions or inventions. The U.S. Copyright Office (USCO) and the U.S. Patent and Trademark Office (USPTO) have unambiguously advised the public that human authorship and inventorship are foundational prerequisites of copyrightability and patentability, respectively. Consequently, a company cannot rely on standard statutory protections to secure exclusive rights in AI-generated assets. Nor can an organization assume that boilerplate SaaS agreement terms will adequately shield its proprietary trade secrets from exposure and destruction.

Because statutory IP protections are unavailable for purely AI-generated outputs, counsel and IP owners face a new paradigm, namely that commercial exclusivity and legal protection must be manufactured through precise, bespoke contract drafting. Counsel must construct contractual guardrails that define the various categories of data relevant to deploying and developing GenAI systems, broadly prohibit AI model training and improvements with customer data,

and protect IP rights after the relationship with the AI vendor ends.

This article provides attorneys with a practical roadmap to protect IP when negotiating AI deployment and development agreements. Specifically, it examines:

- The Emerging AI Tech Stack: An overview of GenAI technology, the concepts, and key terms covered by GenAI agreements and the importance of precise contract definitions.
- Statutory IP Limits vs. Trade Secrets: Analyzing recent court decisions and agency guidance to demonstrate why trade secrets, coupled with contract law, form the primary shield for AI-driven innovation.
- Deployment vs. Development Strategy: Mapping the strategic operational divide between licensing off-the-shelf deployment systems and constructing custom, fine-tuned models, and the granular contractual mechanics required to protect corporate assets.
- Advanced Contractual Mechanics: Providing actionable drafting frameworks to close model-training loopholes, secure customer data ownership, and enforce technical off-ramping protocols at termination.

Ultimately, in-house counsel cannot rely on legacy IT templates or statutory defaults to protect critical enterprise assets. Navigating this shift requires addressing the unique risks that AI systems pose to a company's proprietary information.

Part I: Establishing a Technically Accurate Definitional Baseline

The Hazard of Overbroad Definitions

Enterprise technology contracts traditionally rely on broad, umbrella definitions for technical terminology. In AI procurement, however, copying statutory definitions, such as the expansive definitions found in the EU AI Act or federal executive orders, creates severe contractual hazards. Broad definitions can inadvertently sweep in standard deterministic software, legacy database search algorithms, and basic statistical models.

This creates immediate friction in license scopes, over-inclusive warranty obligations, and unworkable indemnification triggers. Counsel must draft narrow, operationally precise definitions that reflect the technical architecture being licensed or developed.

Deconstructing the AI Stack for Commercial Contracts

To construct an effective agreement, counsel must first distinguish between deterministic software and probabilistic AI systems. Legacy software operates deterministically; given a specific input, the code follows a fixed logical path to produce a repeatable, predictable output. Software performance warranties and uptime metrics were engineered around this predictability.

By contrast, GenAI architectures operate probabilistically. They calculate mathematical likelihoods across volumes of data and billions of parameters to synthesize outputs. Because identical prompts can yield different results, traditional SaaS warranties guaranteeing uniform functional outcomes fail. Contractual representations must instead focus on architectural standards, data hygiene, and processing constraints rather than deterministic outputs.

Generally, agreements that concern the deployment or development of AI technologies will address intellectual property rights and data liabilities for user prompts or inputs and outputs. Where an AI system is being developed,

fine-tuned, or customized, the components of a GenAI system may need to be individually addressed in the contract. In those situations, the agreement's IP terms may need to allocate IP ownership among the following:

- The Core Engine & Algorithms: The underlying, pre-trained base model architecture and source code.
- The Training Data Corpus: The foundational datasets used to train or pre-train the model, which carry third-party copyright infringement risks.
- Model Weights and Parameters: The numerical values adjusted during training that govern how the model processes information, a valuable trade secret in custom developments.

Building on this foundational stack, counsel should anticipate several advanced technical paradigms currently reshaping commercial AI agreements.

Advanced Paradigms In-House Counsel Must Anticipate

As enterprise adoption matures, contracts must account for architectures that go beyond simple prompt-and-response LLMs. In particular, agreements must distinguish between temporary context augmentation and permanent model modification. Retrieval-Augmented Generation connects an LLM to a segregated internal database, querying client data at runtime to generate grounded responses without altering the core model parameters. Fine-Tuning, by contrast, permanently adjusts the model's weights using client data to customize behavior. Ultimately, the contract's IP framework must clearly distinguish between transient data processing and model development, ensuring client data is never converted into vendor-retained training assets or permanent weight adjustments.

Moreover, AI deployments are rapidly shifting from passive prompt-and-response tools to autonomous agentic AI. These systems operate on a continuous observe, reason, act, and execution loop, taking independent steps to fulfill complex directives. When drafting deployment agreements for agentic tools, contractual provisions must account for this shift toward autonomous decision-making. Counsel must establish explicit operational guardrails, strict human-in-the-loop authorization milestones, and defined liability thresholds for instances where an agent executes unauthorized transactions that could release proprietary information into an external application or database.

Finally, enterprise AI relies heavily on standardized integration frameworks – such as the Model Context Protocol (MCP) – to connect AI systems directly to live corporate databases, local file systems, and the internet. While these connector layers eliminate the friction of custom integrations, they simultaneously introduce critical vectors for data exfiltration and IP infringement. Because an MCP connection grants an AI engine structural access to execute actions across connected systems, agreements must establish rigorous security boundaries with explicit usage restrictions to prevent the unauthorized exposure of non-public, proprietary information.

Part II: The Limits of Statutory IP Protection for AI-Generated Assets

The Statutory Vacuum: The Human Requirement

The fundamental premise of federal intellectual property law is that statutory protection exists to incentivize human creativity and innovation. When enterprise systems rely on AI to generate creative works and innovations, organizations limit their ability to copyright and patent their work product. In copyright law, human authorship remains a non-negotiable bedrock requirement. In *Thaler v. Perlmutter*, the D.C. Circuit affirmed that expressive works created exclusively by AI systems are ineligible for copyright protection.^[1] This was solidified when the U.S. Supreme Court denied certiorari in March of 2026.^[2] Similarly, under USCO guidance, prompt engineering and iterative AI output generation are generally insufficient to obtain copyright protection.

A parallel boundary governs patent law. In *Thaler v. Vidal*, the Court of Appeals for the Federal Circuit held that only natural persons – and not AI systems – can be inventors under the Patent Act. [3] Merely presenting a problem to an AI system or requesting a solution is insufficient to obtain a patent. To preserve patentability, enterprise legal teams must implement protocols requiring researchers and engineers to contemporaneously document how they conceived the inventions, including how natural persons formulated specific hypotheses, defined operative parameters, and developed inventive concepts prior to machine interaction.

In light of these limitations, IP rights and ownership of AI-generated outputs should be addressed by the underlying AI deployment or development agreement.

The Litigation Landscape & Training Fair Use

While the lack of statutory IP protection limits output ownership, the copyright battleground over data used for model training has settled into a distinct, two-tiered framework. In the high-profile lawsuits *Bartz v. Anthropic PBC* [4] and *Kadrey v. Meta Platforms, Inc.*, [5] federal courts established that ingesting copyrighted works for model training via statistical pattern analysis is a fair use. Because LLMs process text to learn linguistic probability rather than replicate human expression, the act of model training enjoys protection from a copyright infringement claim.

While statistical processing may be a fair use, courts have drawn a distinction regarding an AI developer's acquisition and storage of copyrighted materials for use as training data. Obtaining and maintaining copyrighted materials from "pirate libraries" creates copyright infringement exposure. [6] A developer cannot insulate the unlawful acquisition of pirated source files simply because the downstream training purpose is transformative. These rulings should influence a party's negotiation of an AI agreement's provisions relating to indemnification for IP infringement, and the inevitably proposed vendor-friendly exclusions.

This legal framework exposes enterprise buyers to a stark commercial reality. Copyright and patent law – traditional protectors of creativity and innovation – provide virtually no shield for generative AI investments. Because machine-generated outputs cannot be copyrighted or patented, competitors can legally copy, redistribute, and monetize an enterprise's unshielded AI-generated innovations without fear of infringement.

Simultaneously, upstream training ambiguities expose vendors and their corporate customers to third-party infringement actions if the underlying models were trained on pirated repositories or if copyrighted materials end up in a customer's output. For these reasons, trade secrets and enforceable contract mechanics have emerged as key means of protecting AI-generated outputs.

The AI and Trade Secret Symbiosis

Because copyright and patent protection are largely unavailable for GenAI output, trade secret law has emerged as the most natural legal framework for securing corporate AI investments. Unlike copyright or patent, trade secret law has no human-authorship or inventorship requirements. Any proprietary information—from raw data to synthetic outputs—can qualify as a trade secret, provided it derives independent economic value from not being generally known and is subject to reasonable efforts to maintain its secrecy.

This creates a dual protection strategy for commercial enterprises. At the system level, corporate developers can protect bespoke model architectures, curated training datasets, and custom-tuned model weights as core trade secret assets. At the output level, corporate licensees can safeguard valuable AI-generated deliverables, such as proprietary software code, complex analytics, or strategic draft documents, under trade secret principles, compensating for the lack of copyright availability.

However, relying on trade secrets exposes enterprise buyers to a critical, potentially catastrophic operational trap. Specifically, losing trade secret status through improper vendor data practices. If an agreement permits the AI vendor to retain, review, or ingest customer inputs or outputs into training datasets, the information's confidential status is likely destroyed. Once proprietary inputs or generated outputs leak into an external model's general knowledge base, secrecy is irrevocably lost, extinguishing trade secret protection.

Part III: The Front-End Strategic Deployment vs. Development Divide

The Two Tracks of Commercial AI Transactions

Presently, there is no standard form of agreement used for the deployment or development of AI technologies. Therefore, counsel cannot approach AI agreements with a boilerplate framework. Instead, drafting and negotiation strategy must be divided into two distinct operational tracks that depend on whether the primary commercial objective is the deployment or the development of an AI system. Deployment agreements involve procuring existing, off-the-shelf AI systems or SaaS platforms that offer AI features. Development agreements, by contrast, govern the creation, fine-tuning, or customization of AI architectures. While IP ownership and usage rights may be negotiated as part of an AI deployment agreement, contracts relating to the development of AI systems require intellectual property allocations across the technical stack.

The Operational Matrix: A Comparative Framework

To select the appropriate contractual framework, in-house counsel can evaluate four primary operational dimensions across the transaction lifecycle:

Operational Dimension	AI Deployment Agreement	AI Development Agreement
Core Licensing Model	Software License / Non-exclusive SaaS subscription	Professional services & software development agreement
IP Ownership Focus	Strict allocation of client Inputs and generated Outputs	Bifurcation of Background IP vs. Foreground IP
Customization Complexity	Low (Off-the-shelf enterprise models or standard RAG)	High (Bespoke fine-tuning, custom model weights, tailored RAG pipelines)
Off-Ramping Priority	Certified data purge; strict prohibition on vendor retention of IP	Verifiable erasure of custom weights, parameters, and specialized training libraries

While the IP focus in AI deployment contracts is typically customer ownership of prompts and output, paired with robust prohibitions addressing whether and how vendors can use and retain “Customer Data,” agreements to develop AI systems demand tailored provisions that treat fine-tuned weights, specialized training datasets, and modified parameters as valuable trade secrets.

Developing the Use Case & Dataset Lineage

It is an AI governance best practice for organizations to formulate AI use cases and find the right tool for the job before entering into an agreement to deploy or develop an AI system. Regardless of whether an enterprise is deploying an existing tool or commissioning a custom model, contract negotiation must be preceded by a rigorous evaluation of the AI model, including dataset lineage. AI models trained on data that contains third-party IP – such as

copyrighted materials scraped from the Internet – create customer exposure to infringement claims. Particularly in custom developments, counsel should require vendors to disclose the lineage of foundational base models and confirm that training sets were legally obtained without reliance on illicit repositories. Notably, California and New York have passed laws requiring companies that fine-tune AI models to publicly disclose high-level summaries of their training data.

Part IV: Deep Dive into Contractual Mechanics & Advanced Drafting

Refining Definitions to Protect Customer IP

In traditional SaaS agreements, Customer Data is typically defined through a static, one-way pipeline model. Its scope is intentionally framed around what the customer uploads to the vendor's cloud infrastructure to be stored, hosted, or processed by deterministic software. In AI contracts, however, contractual IP protections should map to the technical flow of information through the system architecture, including access to data pipelines, websites, and external databases for real-time processing. To construct a resilient legal framework, counsel should establish three definitional pillars that complement each other.

First, the definition of “Prompts” or “Inputs” should cover more than written user prompts entered into a chat interface. Specifically, it should capture all information, files, instructions, and attachments, along with their associated metadata (e.g., document redlines and comments), submitted to or accessed by the tool. To account for modern RAG architectures, MCP connectors, and automated data pipelines, this definition should include enterprise information processed by the AI engine from connected internal databases or external endpoints.

Second, while it is generally standard for customers to own AI system “Outputs,” counsel must carefully define this term to block vendor attempts to claim ownership of synthesized learnings as the mere byproduct of AI system use. Vendors frequently attempt to narrow Output definitions to verbatim text responses, which leaves summaries, secondary analytics, and structural insights unassigned or owned by them. To close this loophole, the contract should define Outputs broadly to encompass all responses generated by the AI system in response to Prompts and all related analyses and derivative innovations.

Finally, to enforce these protections across the commercial relationship, the contract can consolidate both Inputs and Outputs into a single protective umbrella defined as “Customer Data.” Importantly, many vendors are rolling out AI features with their traditional SaaS platforms without addressing the unique risks and nature of AI in the agreement, as well as adding AI features to legacy platforms in the middle of the contract term. Counsel should prepare a template AI Addendum for these situations and ensure that the definition of “Customer Data” aligns with the underlying SaaS Services Agreement's terms. This structural consistency guarantees that the contract's IP protections attach to every phase of the AI processing loop.

Manufacturing the Ownership Framework

Navigating output ownership in AI agreements requires counsel to bridge a distinct gap between contract law and intellectual property jurisprudence. As discussed above, AI-generated outputs are neither copyrightable nor patentable. To establish the appropriate IP protections, counsel should draft an explicit ownership framework that secures every present and future right that can attach to the generated assets. Specifically, the agreement should include an assignment by the vendor to the customer of any rights that may exist in AI-generated output. Such provisions are common and often state, in sum or substance: “To the maximum extent permitted by applicable law, Customer owns all right, title, and interest in and to the Outputs. Vendor hereby irrevocably and unconditionally assigns to Customer all of Vendor's right, title, and interest, including all intellectual property rights, in and to the Outputs.”

This assignment structure prevents the vendor from downstream monetization, re-licensing, or incorporating customer-generated outputs into its own proprietary products. By tethering the assignment to the “maximum extent permitted by applicable law,” counsel captures all existing statutory protections today while automatically absorbing any future legislative or judicial expansions of AI output copyrightability or patentability.

The “Improve the Services” Trap and Closing the Model Training Loophole

In standard enterprise AI agreements, customers will frequently negotiate a high-level representation that the vendor will not “train its models on Customer Data.” While this appears to be sufficient on its face, standard vendor templates routinely undermine this commitment elsewhere in the agreement, often through a broad provision that allows vendors to use customer data to “improve the Services.”

Vendors frequently insert exceptions permitting them to extract, aggregate, de-identify, or anonymize Customer Data to “operate, maintain, enhance, and improve the Services.” In traditional SaaS environments, this language carries minimal risk because it typically covers operational telemetry, bug fixes, and system monitoring. In an AI ecosystem, however, a broad “service improvement” carve-out creates a massive backdoor. Vendors can leverage it to use the information in a prompt, as well as underlying prompt patterns and semantic structures, by asserting that such activities are service improvements rather than formal model training. Without explicit contractual alignment, a covenant against training on customer data becomes effectively meaningless if the vendor retains the right to use that information under the guise of service improvement.

The Solution: A Comprehensive Definition of “Training”

While prohibiting “training” seems adequate, an unqualified ban misses the mark because standard machine learning updates extend beyond the initial training phase. In practice, vendors refine AI architectures through other processes that exploit customer assets without falling under a technical interpretation of “training.” To effectively close this loophole, counsel should avoid relying on a colloquial or undefined use of “training.” Instead, the contract should establish a broad definition of “Training” that encompasses the entire modern machine learning lifecycle, from initial model architecture through human alignment and continuous update loops.

In this respect, a definition of “Training” should capture training and model development, as well as improvements, refinements, optimizations, and specific enhancement processes, such as fine-tuning, reinforcement learning, and transfer learning. Paired with such a definition, customers should negotiate for the vendor’s contractual promise not to use Customer Data, Inputs, and Outputs for the Training of “any” AI model – so as to cover third-party models – but “solely to provide the Services for Customer.” By extending this protection to the actions of the vendor’s third-party sub-processors, customers can put themselves in a position to avoid one of the most common pitfalls in modern enterprise AI agreements and maintain control over their proprietary information.

Non-Content AI Data Ownership and Usage Considerations

When negotiating AI agreements, counsel typically focuses on protecting “Customer Data.” However, AI use generates an array of data in the form of operational artifacts that vendors commonly seek to own and use for any purpose, such as:

- System metadata: Document lengths, file schemas, and token consumption volumes.
- Behavioral telemetry: User query frequency, edit patterns, and latency thresholds.

- Usage data: Workflow sequences, application integrations, and tool-calling triggers.

Vendors routinely carve out non-content data from the definition of “Customer Data,” classifying it as vendor-owned “AI Usage Data,” or “AI Performance Data.” The trap lies in the nature of modern machine learning architectures, because models do not require raw text to extract high-value corporate intelligence. Rather, AI systems learn from patterns of data.

By analyzing metadata patterns, structural embeddings, and behavioral dynamics over time, a vendor can effectively reconstruct user profiles and gain insight into their usage patterns, behaviors, and preferences. When vendors retain comprehensive rights to analyze and build internal tools with this non-content data, they can monetize the customer’s operational footprint.

To eliminate this vulnerability without disrupting basic platform functionality, counsel should focus on establishing tightly delineated, mutually exclusive contract definitions for AI Usage Data and AI Performance Data. Vendor templates often blur the line between operational metrics and customer assets. Allowing “Usage Data” to serve as an umbrella term that encompasses underlying prompt metadata, schema structures, and vector embeddings runs the risk of the customer forfeiting the proprietary information inherent in this information. Specifically, vector embeddings preserve the core semantic logic of proprietary documents, and custom schemas embody trade secret system architectures.

To protect proprietary information and prevent downstream model training, counsel should carefully negotiate data category definitions and strict boundaries on ownership and usage. For example, by defining “AI Performance Data” solely as the high-level, technical telemetry required to maintain system stability, measure network throughput, and calculate billing units, customers avoid opening backdoors for vendors to capture, analyze, or train on the structural and semantic footprints of their IP. Importantly, definitions of non-content data should explicitly exclude Inputs, Outputs, and Customer Data; and, by broadly defining “Outputs” as noted above, customers can maintain control over the data derived from or reflective of the proprietary information explicitly generated by the AI system.

Part V: Advanced Development Dynamics & Technical Safeguards

Custom AI development demands a shift toward active IP engineering. As enterprises move from off-the-shelf AI applications to bespoke fine-tuning, RAG architectures, and proprietary model builds, contractual mechanics must pivot accordingly. Negotiating AI development agreements requires counsel to anticipate how curated datasets, parameter adjustments, and fine-tuning artifacts interact with IP concerns. The fundamental objective shifts from protecting inputs and outputs to securing the IP created during the development lifecycle.

Navigating the Customization Spectrum: Three Negotiation Paths for Model Ownership

When an organization contracts to develop an AI model, the IP focus evolves to include ownership rights in the data used to create the system and the custom AI deliverables themselves. In practice, counsel must navigate three distinct ownership structures, with each carrying different risk profiles and commercial outcomes.

The first structure is exclusive customer ownership. The best outcome for enterprise risk mitigation is full, unencumbered customer ownership of all AI models and adjustments developed for them, such as fine-tuned model weights. Effective IP protection requires counsel to understand the nature of the AI technology being developed. For example, RAG architectures involve the selection of curated data libraries and MCP connectors may permit access to proprietary or external databases. For this reason, counsel should ensure that the contract gives them ownership of not only the custom model, but the data pipelines, index schemas, MCP server integrations, and curated knowledge

bases created to support the system. Securing exclusive ownership guarantees that the developer cannot leverage the customer's domain expertise or trade secrets to build a competing model for third parties.

The second scenario is developer ownership with an exclusive license back to the customer. Where the developer provides a deeply entrenched foundational model or proprietary core architecture, the vendor will strongly resist transferring outright ownership of fine-tuned weights. A common compromise is developer ownership paired with a highly restrictive, exclusive commercial license back to the customer.

To adequately protect company IP and competitive positioning under this framework, counsel should strictly negotiate the parameters of exclusivity. Specifically, the license should be more than a basic grant to use the software. Here, customers should seek to establish an exclusive field-of-use or industry-specific commercial silo. The contract should explicitly prohibit the developer from deploying, re-licensing, or incorporating deliverables and functionally equivalent configurations derived from the customer's training data into products offered to competitors.

The last scenario is joint ownership of custom AI deliverables, which is almost always the worst commercial outcome because it creates operational gridlock and erodes competitive value. For example, the vendor-joint IP owner may be able to license the technology to the customer-joint IP owner's competitors without consent or royalty-sharing (unless modified by contract). Additionally, enforcing jointly owned IP against third-party infringers may require both owners to join the lawsuit. If a rival misappropriates the model architecture, the developer may refuse to litigate against its other clients or partners, effectively paralyzing the customer's legal remedies. Also, maintaining trade secret status requires reasonable efforts to preserve secrecy. Granting a developer co-ownership over a model embodying customer know-how gives a trade secret defendant an argument that the customer failed to use reasonable measures. Unless the agreement incorporates cross-licensing restrictions, profit-sharing mechanics, and enforcement protocols, customers should avoid joint ownership in favor of clear sole ownership coupled with defined vendor data usage and retention restrictions.

Technological Segregation Of Customer Data

Protecting enterprise IP in custom development environments involves enforcing technical boundaries. In multi-tenant cloud architectures or shared model hosting environments, logical or physical data segregation is critical to ensuring that customer datasets, prompt embeddings, and fine-tuned parameters do not intermingle with the vendor's base models or another customer's workflows.

Agreements pertaining to the development and use of custom AI models should include technical specifications that preserve the separation of the customer's IP, such as:

- **Tenant-Level Isolation:** Storing customer vector databases, RAG indices, and custom weights in dedicated, encrypted partitions that prevent cross-tenant data leakage or inadvertent extraction by secondary model processes.
- **Derived-Asset Containment:** Ensuring that all temporary caches, intermediate embeddings, and transient execution files generated during model inferencing or fine-tuning are dynamically purged or strictly confined within the customer's tenant boundary.
- **Direct Derivation Warranties:** Negotiating for warranties that customized outputs and model enhancements are derived exclusively from the customer's secure instance and authorized inputs, rather than from shared data lakes.

Enforcing these technical safeguards guarantees that the customer's proprietary assets remain strictly within its

control, mitigating the risk of trade secret exposure through shared infrastructure and datasets.

Downstream Protection and Mitigating Infringement Risk

Custom AI architectures rarely rely solely on internal enterprise data. If the developer intends to use datasets that carry restrictive license terms or unverified intellectual property rights, the custom model can expose the customer to IP infringement claims. Thus, counsel should require representations that the developer holds all necessary rights and licenses for any third-party data or weights integrated into the build. Relatedly, the development agreement should include warranties that no components were used in a manner that would force the licensing of the enterprise's proprietary fine-tuned weights. Finally, counsel should ensure that IP indemnification expressly flows through to cover third-party claims arising from underlying training datasets and integrated model components.

Part VI: The Off-Ramping Protocol

The Technical Barrier: The Impossibility of “Unlearning”

Standard SaaS off-boarding relies on a simple premise; that is, upon contract termination, the vendor deletes the customer's tenant database and returns or purges its raw data. In generative AI environments, however, standard data destruction protocols completely collapse against the reality of neural networks. Once prompts, fine-tuning datasets, and other customer data are ingested during a training or optimization run, that data is compressed into billions of floating-point numbers. A trained neural network does not store data like a database. Instead, it absorbs statistical patterns and relationships. Consequently, a trained model cannot simply “unlearn” specific data points without retraining the entire model from scratch.

This technical reality reinforces why front-end, preventative “no training” covenants are so vital. Once customer data crosses the threshold into a vendor's data training pipeline, a simple post-termination data deletion request cannot undo the weight adjustments made with an organization's IP.

A Deletion Framework for Custom Models

Traditional boilerplate provisions that give vendors 30 days to return or destroy customer data fail in custom AI deployments. These generic clauses leave vendors with the ability to retain derived model weights, vector indexers, cached embeddings, and temporary fine-tuning artifacts under the guise that they are internal operational assets rather than “Customer Data.” To prevent trade secret leakage, counsel should draft post-termination obligations that encompass all secondary, AI-derived artifacts.

When custom models or RAG databases are built with enterprise data, the agreement should mandate technical destruction mechanisms:

- **Derivative Weight Destruction:** Wiping all fine-tuned model checkpoints, files, logs, and custom training datasets.
- **Vector Store Purging:** Complete deletion of vector database collections, semantic indexers, and document chunking caches.
- **Cryptographic Erasure:** This renders residual model weights and embeddings permanently unrecoverable by the vendor.

To ensure compliance, counsel should require written confirmation of destruction from an authorized vendor

representative upon completion, backed by the right to request reasonable technical verification or system logs demonstrating that all residual weights, indices, and training artifacts have been permanently purged.

Conclusion and Key Strategic Takeaways

Navigating enterprise AI transactions requires counsel to look beyond standard software licensing concepts and actively engage with the underlying technical architecture. Protecting enterprise value in the age of AI demands an IP-first approach that addresses how data flows through, alters, and persists within AI systems. In so doing, counsel must look past static data storage and account for the fluid nature of AI systems. Without targeted definitions and drafting that closes a vendor's model-training and data retention loopholes, enterprise domain expertise risks being silently harvested to train commercial tools for market rivals, putting a company's valuable IP in jeopardy.

Footnotes

- ¹ 130 F. 4th 1039 (D.C. Cir. 2025), cert. denied, 146 S. Ct. 1777 (2026).
- ² *Thaler v. Perlmutter*, 146 S. Ct. 1777 (2026) (mem.).
- ³ 43 F.4th 1207 (Fed. Cir. 2022), cert. denied, 143 S. Ct. 1783 (2023). Under the USPTO's current guidance, there is no separate or heightened inventorship standard for AI-assisted inventions. AI systems are treated as tools, and inventorship turns on whether a natural person made the requisite contribution to conception under the traditional standard. See *Revised Inventorship Guidance for AI-Assisted Inventions*, 90 Fed. Reg. 54636 (Nov. 28, 2025).
- ⁴ 787 F. Supp. 3d 1007 (N.D. Cal. 2025).
- ⁵ 788 F. Supp. 3d 1026 (N.D. Cal. 2025).
- ⁶ *Bartz*, 787 F. Supp. 3d at 1025-27, 1033.

[https://prod.resource.cch.com/resource/scion/document/default/\(%40%40PAA01%20WK-JSTORY20261001-3\)paa01f16bd8a5413046e88082315bc5810bc8?cfu=Legal&cpid=WKUS-Legal-Cheetah&uAppCtx=cheetah](https://prod.resource.cch.com/resource/scion/document/default/(%40%40PAA01%20WK-JSTORY20261001-3)paa01f16bd8a5413046e88082315bc5810bc8?cfu=Legal&cpid=WKUS-Legal-Cheetah&uAppCtx=cheetah)