BCLP. Client Intelligent

Insights

AI CONTRACTS: WHOSE PROMPTS, WHOSE OUTPUTS, WHOSE RIGHTS?

Nov 17, 2026

Summary

As businesses increasingly deploy generative AI tools (both for internal business purposes and to support delivery of products and services to customers), questions arise about ownership and usage rights in both the outputs and the prompts used to generate them. Contracts typically seek to address these questions by allocating IP ownership and usage rights, but uncertainty remains about the protection these provisions offer in practice, given courts and legislatures worldwide are still grappling with the issue of IP ownership in AI-generated outputs. This article explores the considerations for customers and vendors when negotiating rights in outputs and prompts under AI contracts as the law races to catch up with the practical uses being made of AI across every industry and sector and increased usage of AI agents.

Assignment and ownership of outputs

The customer perspective

Current copyright law does not protect works generated entirely by AI without human involvement that meets the threshold for copyright protection. Human users may obtain copyright protection where they make their own creative contribution when using AI tools, but the threshold for subsistence of copyright protection remains legally unsettled.

Generation of outputs by an AI system will not generally give the vendor any IP rights in such outputs, given that its contribution is typically limited to the provision of the underlying technology. To the extent that protectable IP rights exist, they are more likely to vest in the customer (or its personnel) if their contribution meets the threshold for copyright protection in the relevant jurisdiction where the AI tool is used. This uncertainty is the principal reason for including assignment and license clauses in AI contracts.

© 2026 Bryan Cave Leighton Paisner LLP.

Where a contract provides for the vendor to assign or license IP rights in AI-generated outputs to the customer, customers should scrutinize which rights the provision actually transfers or grants. The practical effect of such provisions is limited by the fundamental principle that a party can assign only rights it owns and grant licenses only within the scope of its ownership or authority.

The composition of the underlying model training data compounds this limitation. Training data typically includes licensed and publicly available material from third-party sources that the vendor does not own. If that material is reproduced, in whole or in part, in outputs, an assignment or license between the vendor and customer is unlikely to confer clean title or legitimate usage rights in those outputs. The clause may therefore fail to achieve its intended effect.

Before licensing an AI tool, customers intending to commercialize its outputs should, to the extent possible, conduct due diligence on the system's training data to assess whether ownership of IP rights in those outputs is realistically attainable. This is particularly important where customers upload their own proprietary material and expect to retain control over its use and any resulting outputs.

Alongside this due diligence, customers should seek broad contractual rights to reproduce, distribute, modify and commercially exploit outputs, supported by appropriate non-infringement warranties and indemnities against third-party claims. These provisions offer practical protection irrespective of how the underlying ownership position is ultimately characterized.

The vendor perspective

Vendors should be cautious when agreeing to assign rights in generated outputs, which may incorporate data or content owned by third parties. Where outputs are generated from training data comprising licensed, proprietary and public-domain material, vendors will typically be unable to warrant clean title. By agreeing to a blanket assignment of all content within outputs, a vendor risks making an unqualified commitment to assign rights it cannot convey.

Vendors may be better served by allowing ownership to be determined by operation of law rather than purporting to assign rights they do not own. Under this approach, any IP rights arising in AI-generated outputs, to the extent protectable, may vest in customers only insofar as the outputs are protectable and the relevant material in those outputs is not already owned by a third party.

Where a vendor's own IP forms part of the system's training data, the vendor should seek express carve-outs reserving all rights in that IP. In particular, the contract should exclude the vendor's proprietary materials from any rights that may vest in the customer in generated outputs. Without these safeguards, a vendor risks inadvertently transferring or relinquishing its own IP rights.

Vendors should also be cautious when providing non-infringement warranties or indemnities against third-party IP claims relating to generated outputs. If training data includes third-party content licensed on terms that prohibit sublicensing or further usage rights for downstream

© 2026 Bryan Cave Leighton Paisner LLP.

recipients, the vendor may be unable to warrant that outputs are free from third-party encumbrances. Such warranties or indemnities may expose the vendor to liability it cannot adequately manage or pass through to the relevant rights holders. Vendors should therefore qualify these provisions to reflect the limitations of the underlying third-party licenses.

Prompts: usage rights and competing interests

A related but distinct issue concerns the prompts customers submit to generate outputs. Prompts frequently contain commercially sensitive material, including strategic plans, pricing information, draft work product and proprietary data. Customers will generally wish to keep this material confidential and restrict the vendor's ability to use it.

Vendors, by contrast, have a commercial interest in using customer prompts to improve their models. A vendor seeking permission to reuse prompts for model training will need to ensure that any confidentiality commitments to customers accommodate such use.

Some vendors expressly disclose that customer-generated prompts are not treated as confidential and may be used to train their models. Even without such disclosure, customers should not assume that silence on this point in a vendor's standard terms operates in their favor. It also remains unclear who owns the rights to prompts that a customer's employees generate or modify using AI tools, an increasingly common practice that may reduce the human contribution below the threshold required for copyright to subsist.

The treatment of prompts has therefore become a key negotiating point. Customers increasingly seek broad no-training covenants prohibiting vendors from using customer data contained in prompts to train their models. It is also a live issue when negotiating non-disclosure agreements, where a party disclosing confidential information will want to ensure their information is not exposed to an AI model which would allow that information to be used to train the model.

A common compromise between these competing commercial interests is to distinguish customer data contained in prompts and outputs from purely operational usage data. The parties may agree to prohibit training on customer content while permitting use of operational data for purposes such as improvement of services.

An evolving landscape

As AI systems advance rapidly, IP ownership in outputs and usage rights in prompts will remain fast-evolving and hotly negotiated issues in AI contracts. The protections that matter most in practice will vary from case to case, requiring careful, tailored drafting and an understanding of what is going on 'under the hood' of a specific AI tool, rather than reliance on standard positions.

© 2026 Bryan Cave Leighton Paisner LLP.

Related capabilities

- Intellectual Property & Technology Disputes
- Digital Transformation & Emerging Technology

Meet the team



Vishnu Rege

Associate, London

vishnu.rege@bclplaw.com
+44 (0) 20 3400 4552



Anna Blest

Counsel, London

anna.blest@bclplaw.com
+44 (0) 20 3400 4475

This material is not comprehensive, is for informational purposes only, and is not legal advice. Your use or receipt of this material does not create an attorney-client relationship between us. If you require legal advice, you should consult an attorney regarding your particular circumstances. The choice of a lawyer is an important decision and should not be based solely upon advertisements. This material may be “Attorney Advertising” under the ethics and professional rules of certain jurisdictions. For advertising purposes, St. Louis, Missouri, is designated BCLP’s principal office and Kathrine Dixon (kathrine.dixon@bclplaw.com) as the responsible attorney.

© 2026 Bryan Cave Leighton Paisner LLP.