

# Does Your AI Agent Need Its Own Software License?

**New Media and Technology Law Blog** on August 24, 2026

- AI agents create contracting and pricing issues for both traditional licensed software and SaaS by weakening the assumptions behind per-seat and other usage-based pricing arrangements.
- As organizations deploy agents that can act autonomously and generate substantial usage, the parties to software license and SaaS agreements will have to address how agent access should be authorized, measured and priced when the traditional “per-seat” metrics are challenged by these new ways to work.
- The issue spans both traditional software licenses and SaaS arrangements. For licensed software, AI agents may raise questions about the scope of licensed deployment and use; in SaaS arrangements (where the customer typically receives access rights rather than a license to the underlying software), agents can challenge user-based subscription, access and consumption entitlements.

In a recent ruling, the Federal Circuit considered the Court of Federal Claims’ award of copyright damages to software company 4DD Holdings, LLC (“4DD”) and its subsidiary in a long-running litigation brought against the federal government, with the court holding that case law does not require automatic adoption of royalty rates set out in licensing agreements when calculating infringement damages, but vacating portions of the award and remanding based on errors in the damages analysis. ([4DD Holdings, LLC v. U.S.](#), No. 2024-1996 (Fed. Cir. July 16, 2026)). In its 2023 [decision](#), the trial court found that the government (through its contractor) deployed the software beyond the licensed number of cores and seats and had exceeded the scope of its licenses for 4DD’s software by creating thousands of unauthorized copies, including backup and testing copies created through virtual machine cloning, even though the license permitted only a single backup copy and authorized deployment to a specified number of seats and processor cores. The trial court treated those out-of-scope reproductions and deployments as copyright infringement and awarded 4DD \$12.7 million.

Beyond the important holdings in this case about the calculation of damages in a software license case, it was other aspects of this dispute and the software license at issue that sparked our interest. The *4DD* dispute did not necessarily concern the treatment of AI agents under a software license, but the case illustrated what could happen when technical deployment practices exceed contractual limits. Specifically, the dispute brings to mind an emerging issue about how traditional software and SaaS agreements will have to (and already have begun to) evolve in the era of AI agents, which are beginning to put pressure on the classic “per seat” pricing or entitlement framework. Whether enterprise software is licensed on premises by customers or accessed remotely in a SaaS model, the use of AI agents creates challenges in the customer-provider relationship, but the nature of the challenge depends on the pricing model: agents put pressure on per-seat SaaS economics while also complicating traditional software licensing models. Future disputes involving AI agents exceeding a contractual “seat” entitlement may create a contract/licensing issue (but not necessarily rise to the level of a copyright claim involving the unauthorized reproduction of software, as was the case in *4DD*).

In short, AI agents break the traditional assumption that identity, access, activity and value are all measured by the same licensing unit—and current vendors are already responding in different ways.

### **Per-seat licensing & AI agents**

While it is true that no two human users will have the same usage rate for software, the traditional headcount-based SaaS pricing and subscription model was nevertheless workable because of the practical constraints of human usage. Seat-based pricing implicitly tolerates differing levels of activity among human users because the licensed unit remains the person rather than each discrete action. AI agents, by contrast, may either generate a volume of activity that previously would have been associated with multiple human users, or they may be narrowly configured to perform a limited number of tasks. Thus, their deployment is exposing ambiguities in traditional enterprise software arrangements because the language of the agreement often does not specify how autonomous software users should be counted. For example, seat-based licensing raises questions about who or what counts as an authorized user while core- and instance-based licensing raise questions about how agent-generated workloads, virtualization and ephemeral resources are measured.

The issue of AI agents can create contracting and pricing issues for both traditional licensed software and SaaS arrangements.

With traditional licensed software, the customer generally receives contractual/copyright rights to install, reproduce or use software within defined limits. An AI agent can complicate those limits by increasing automated use, interacting through service accounts, or otherwise indirectly accessing software in ways the agreement did not contemplate. This is where concepts such as authorized users, instances and indirect access matter. At a basic level, the questions here are more likely to involve what the customer is contractually permitted to do with the software.

With SaaS, the customer typically does not receive a copyright license to install the provider's software, but instead receives contractual rights to access and use a hosted service. But AI agents pose an analogous problem for the subscription and pricing model. The central question here is often whether one human seat can support multiple agents generating substantial activity, or whether those agents require separate seats or consumption charges.

At a high level, the long-standing value proposition that more seats equals more productivity for the customer (and scalable revenue for the software vendor) is starting to be squeezed by new technologies. Indeed, for SaaS providers in particular, per-seat pricing presents perhaps the clearest challenge for software providers in the age of AI, where the use of agents by a customer's employees could lead to compressed seat counts, as an enterprise customer that previously needed 100 seats may eventually need only 50 employee-users supervising AI agents. The vendor would therefore receive less revenue even though its software would likely process the same or greater volume of activity. In response, vendors could, in certain instances, characterize agents as additional digital users or "seats" (e.g., one employee supervising ten agents could theoretically require eleven seats rather than one) or else move toward a hybrid, consumption-based structure.

This new reality creates a fundamental question about how enterprise software should be priced in the future: Is an agent simply an extension of a human user, a separate licensed or authorized user, a machine process, or merely a source of metered consumption? The market is experimenting with different approaches to the relationship among human users, non-human identities, and consumption, rather than converging on one post-AI agent pricing framework. Yet, many negotiated enterprise agreements drafted before widespread agent deployment may contain definitions and metrics developed principally for human users and more conventional infrastructure in mind. Thus, going forward, customers and providers should address agent access expressly in the event running AI agents becomes an unexpected compliance event or an updated pricing model transition becomes a source of unbudgeted fees for the customer. This emerging issue makes product-specific definitions and negotiated license terms increasingly important.

### **An increasingly important transactional issue**

The basic problem is that traditional metrics assume a relatively stable relationship between the licensed unit and the customer's use:

- A per-seat license assumes that a human user accesses and derives value from the software.
- A per-core license assumes that processing capacity is a reasonable proxy for use.
- An instance- or server-based license assumes that deployments are relatively identifiable and persistent.

### **AI agents may present certain ambiguities under existing agreements**

These include:

- Does a "user" have to be a natural person?
- Does each agent require a separate license, seat, subscription entitlement, or other access right? Does an agent authenticate under a separately provisioned identity or act under authority delegated by a licensed or authorized user, and how does the underlying agreement treat each model?
- May one licensed user use an unlimited number of agents?
- Is unattended or autonomous access treated differently from human-interactive use?

- Does an agent that accesses several applications require a credentials for each?
- Are agent API calls and agent-to-agent communications licensed or authorized uses?
- Does spawning temporary or parallel agent instances trigger additional user, instance, API or consumption entitlements?
- Can an agent use the software on behalf of multiple employees, affiliates, customers or business units?
- Does the license or SaaS agreement prohibit credential sharing even where the “sharing” occurs between a user and the user’s AI agent?

Many legacy definitions of “authorized user,” “access,” “use,” “device,” or “instance” were drafted before autonomous AI agents became more common. A licensor or provider may argue that every separately credentialed agent is another user, while a customer may counter that the agent is merely a tool acting on behalf of an already purchased “seat” or employee. Unless the contract addresses that distinction, both sides may face uncertainty.

This uncertainty, however, is not entirely new. Enterprise software licenses have long addressed “indirect access” or “multiplexing,” in which individuals receive the benefit of licensed technology through middleware, rather than signing into the licensed product themselves. In many cases, a software vendor will treat certain users who indirectly trigger and receive value from automated flows as users requiring licenses. Multiplexing is an analogue, not necessarily the answer to agent licensing, as an autonomous agent may act as its own non-human principal, act under delegated authority, or on behalf of multiple users.

The framing of the problem itself may also contribute to the uncertainty. Asking only whether an agent is a “user” collapses several questions that should be considered separately: whether the agent is authorized to access the software; whose authority and permissions it uses; on whose behalf it acts; and how its activity should be measured and priced. Some emerging agent-identity systems (e.g., Microsoft Entra Agent ID) are being designed to give agents distinct identities, limited permissions, identifiable human sponsors, and auditable activity. Ultimately, a contractual classification of agent use should be paired with operational controls (e.g., identity records, usage logs and true-up rights) once a framework is negotiated rather than relying on expanding or tailoring the definition of “user” alone.

For contracting purposes, companies may need to distinguish between the rights required to deploy an AI agent and the licenses, subscriptions, and access rights required for the applications, data and services that the agent accesses.

## **Practical provisions to address**

Below are some issues for organizations to consider as the nature of work changes to include greater support from advanced AI agents and tools.

For licensees or customers:

- a definition of “authorized user” and what rights a user’s AI agents may have to use APIs and create accounts;
- confirmation that an agent acting exclusively for a licensed or authorized user does not require an additional seat;
- caps or predictable pricing for agent usage;
- clarity on whether, and under what circumstances, seat-based charges and consumption charges apply cumulatively;
- reasonable cure and true-up mechanisms before breach or termination;
- transparency into how the vendor measures agent activity.

For licensors or providers:

- whether each agent or service identity requires authorization;
- limits on the number of agents associated with a human user;
- usage by agents serving multiple users or external customers;
- metering of API calls, transactions, tokens, workflows, or completed outcomes;
- audit rights covering non-human identities and automated access;
- pricing or true-up mechanisms applicable when agent usage materially increases system consumption while reducing human seat counts.

### [Related Professionals](#)

---

- **Wai L. Choy**  
Partner