

MAYER | BROWN

AI SUMMIT 2026

New York

Thought Leadership Materials

Panel Session 3: Contracting for AI

FEBRUARY 17, 2026

CONTRACTING FOR AGENTIC AI SOLUTIONS: SHIFTING THE MODEL FROM SAAS TO SERVICES

AUTHORS: ROHITH P. GEORGE, JOE PENNELL, BRAD L. PETERSON, OLIVER YAROS

AT A GLANCE

As agentic AI products shift from passive tools to autonomous actors, we see a move beyond traditional SaaS contracting models to a hybrid approach incorporating BPO-style clauses, including clauses covering service definitions, warranties, outcome-based SLAs, broader indemnification, governance and audit rights, and data ownership.

For years, contracting for generative AI (GenAI) products has largely settled into a familiar Software-as-a-Service (SaaS) model: the provider makes the GenAI product available on its platform and the customer company is responsible for how it is used. This model often makes sense when the AI product is a passive tool—a co-pilot that suggests but does not act.

Agentic AI, however, does not neatly fit into this contracting model. “Agentic AI” refers to systems that can autonomously plan and execute multi-step tasks to achieve a goal. Instead of just suggesting content, they are autonomously taking action on the company’s behalf.

Agentic AI includes a spectrum of products. At one end of the spectrum, there are general-purpose agentic AI tools that allow companies to develop and build their own AI agents. The company’s team has substantial ability to train, fine-tune, adapt, program and otherwise direct those AI agents. At the other end of the spectrum, there are agentic AI solutions. Agentic AI solutions are developed by providers to perform specific functions, such as handling payment inquiries from suppliers or helping employees access their benefits. The company’s team has limited or no ability to affect how the AI agents operate. In between, there are varying levels of control and company involvement.

As an agentic AI solution shifts to acting autonomously on a company's behalf, the nature of the provider’s relationship with the company shifts from licensing a tool toward providing a service. With that change in relationship, we see the contracting model shifting from a SaaS contracting model—with limited performance guarantees and software-focused risk allocations—to a more service-oriented contracting model. That service-oriented model would require defining the service, setting guardrails and governance rights and obligations, creating incentives for proper oversight and management on each side, and allocating liability for service failures.

Fortunately, these are not wholly new issues. The business process outsourcing (BPO) industry has established market terms and conditions to address these thorny issues when a company hires a service provider to perform business services using people. The challenge now is to adapt those concepts fairly and appropriately from the BPO industry when hiring a service provider to deliver services using AI agents.

This Legal Update identifies six critical clauses where the standard SaaS contracting framework is a poor fit for agentic AI solutions and proposes updated versions of BPO-style solutions as a more balanced, appropriate starting point for both buy-side and sell-side negotiations.

1. DEFINITIONS & SCOPE OF SERVICE

The SaaS Clause: A standard SaaS agreement defines the “Service” as a hosted software platform, with the company receiving a non-exclusive right to access and use that platform. The provider is responsible for providing the platform; the user is responsible for all actions taken with the platform.

The BPO-Style Solution: The “Service” would be defined as the set of tasks and responsibilities that the provider agrees to complete using AI agents. This definition of services would explicitly define the provider’s “delegation of authority” and any “policy guardrails.” The delegation of authority would outline what the provider can do using AI agents (e.g., offer an on-site service call) and what it cannot do using AI agents (e.g., offer refunds or accept liability for claims). The policy guardrails would specify how the AI agents can operate, including mandatory escalation triggers for human-in-the-loop (HITL) approval. This provides clarity for the company and a defensible liability guardrail for the provider.

PRACTICE TIPS FOR COMPANIES: KEY QUESTIONS TO ASK

- What parts of the *business process* will be performed with AI agents?
- Which steps are critical to define in the policy guardrails?
- What is the precise delegation of authority? What is the exact threshold where the agentic AI solution *must* stop and ask a human for approval?

2. SERVICE WARRANTIES

The SaaS Clause: “THE SERVICE IS PROVIDED AS-IS, WITH ALL FAULTS.” This is the opening of an extensive, all-caps disclaimer in many SaaS agreements for AI products. A SaaS provider may be willing to offer a warranty that its product will perform in material conformance with its documentation. However, SaaS-based AI providers argue that they cannot offer even that warranty given the probabilistic nature of AI.

The BPO-Style Solution: A BPO-style approach would include performance warranties. BPO agreements generally include warranties that services will be performed in a good, professional, diligent, and workmanlike manner in accordance with industry standards. In this solution, the contract would apply that warranty both to the work of the people who create, monitor, and maintain the AI agents and expressly also to the work performed by AI agents as if performed by people. In addition, BPO agreements routinely include warranties of compliance with key restrictions, so this solution might include a warranty that the Services will be in compliance with law and in material conformance with the delegation of authority and policy guardrails as defined in the Agreement.

PRACTICE TIPS FOR COMPANIES:

- Providers will generally not accept warranties of perfection in delivery of services. Quality is thus

addressed by the more circumscribed warranties described above or Service Level Agreements (as described below).

- The warranty of compliance to the delegation of authority and policy guardrails is directly linked to Clause 1. The better you define the delegation of authority and the policy guardrails, the more willing a provider may be to warrant that its AI agents will stay within that defined scope.

3. SERVICE LEVEL AGREEMENTS

The SaaS Clause: Service Level Agreements (SLAs) are technical and measure platform availability. 99.99% “uptime” is a common standard in a SaaS agreement. This provides little comfort if the agent is “up” but making costly errors.

The BPO-Style Solution: SLAs would be operational and measure outcomes rather than solely availability. Service credits can still be a remedy, but they are triggered by performance failures, not just downtime.

Key metrics for an agentic SLA might include:

- **Accuracy:** e.g., 99% of invoices processed correctly against the purchase order.
- **Timeliness:** e.g., 99% of support tickets actioned within the required service window.
- **Satisfaction:** e.g., <1% of autonomous actions lead to consumer complaints.

PRACTICE TIPS FOR COMPANIES:

- Talk to your internal business stakeholders and ask, “What does good look like for this agent?” Translate their business-focused answers (e.g., “it doesn't make mistakes,” “it's fast”) into measurable metrics like “Accuracy” and “Timeliness.”
- Propose these outcome-based SLAs instead of accepting the provider's standard “uptime” SLA.
- Note that this may also be in the interest of the provider, as this will help the provider define the remedies for poor performance (e.g., if the provider's position is that service credits are a sole and exclusive remedy).

4. INDEMNIFICATION

The SaaS Clause: In many SaaS agreements, indemnities are often narrow. In some cases, the provider's only indemnity is to defend and hold harmless the company from third-party IP infringement claims in relation to the Services.

The BPO-Style Solution: Indemnities are often broader, as they are designed to cover risks that arise from the way the service is performed. For agentic AI, a company may seek indemnification from the provider for a broader range of third-party claims arising from the agent's autonomous actions in breach of the delegation of authority or the policy guardrails.

Examples could include claims where the agent discriminates in an automated hiring workflow or breaches the policy guardrails. These indemnities may be balanced with provider-favorable carve-outs (e.g., indemnification would not apply to harms caused by (a) company misconfiguration, (b) faulty company data, or (c) an action the agent escalated and the company's HITL *explicitly approved*).

PRACTICE TIPS FOR COMPANIES:

- Consider proposing a provider indemnity for third-party claims “arising from the agent's autonomous performance of the Services.” Be prepared to narrow this to specific areas of risk.
- To make this reasonable, proactively offer carve-outs for your own company's failures, like providing bad data or a bad decision.

5. GOVERNANCE & AUDIT RIGHTS

The SaaS Clause: Audit rights may be limited to the provision of a multi-customer SOC 1 or SOC 2 audit report, and the provider may even have the right to audit the company for seat-license or usage overages (a legacy from on-premise software licensing models).

The BPO-Style Solution: A BPO-style approach would be to provide broader company audit (or transparency) rights. A company delegating a core function may require a contractual “right to transparency.” For example:

- **Technical:** The provider would have an obligation to maintain decision logs for all decisions by AI agents and customer would have the right to audit any agent's decision logs to help answer the question, “why did it do that?”
- **Operational:** The right to formally assess AI agent performance against the agentic SLAs and other contractual commitments.

PRACTICE TIPS FOR COMPANIES:

- Before signing, conduct a technical and operational analysis to determine how to audit whether the provider’s work complies with requirements, then ask for contractual rights to conduct that audit.
- Require that “decision-logs” and other key records be maintained in structured, legible form so that they are useful.

6. DATA, IP RIGHTS & MODEL TRAINING

The SaaS Clause: SaaS terms can sometimes grant the provider a broad, perpetual license to all data generated by or through the platform (including company inputs and AI outputs) to “improve the service.” This often means the provider can train its models on the company’s confidential data.

The BPO-Style Solution: While BPO providers may similarly ask for rights to use company data, a BPO relationship is a service provider/processor relationship and therefore the lines are much clearer. Under a BPO-style approach:

- **Data & IP Ownership:** The contract would state unambiguously that the company owns (a) all data that is submitted to or obtained by the agent (the inputs), and (b) all outputs created (including IP rights therein) or generated by the agent in the performance of the service. Any license to use such company data (e.g., in original or deidentified form) would be carefully negotiated.
- **Model Training Rights:** The contract would explicitly prohibit the provider from using company data (inputs or outputs) to train its models unless the company consents to such training, perhaps for its own benefit or for collective benefit.

PRACTICE TIPS FOR COMPANIES:

- Find the “Data Use” or “License to Provider” clause and insert an explicit prohibition against using company data to train, fine-tune, or otherwise improve any AI model without your approval. Be

open to discussions about win-win scenarios regarding use of data.

- With some exceptions, ensure that, as between the parties, you own the generated outputs. You are paying for a service that *creates* those outputs on your behalf. However, providers may want you to acknowledge that the output generated for you may be the same or substantially similar to an output generated for another customer.

CONCLUSION: FORGING THE HYBRID CONTRACT

The path forward for procuring agentic AI will likely not be to scrap the SaaS contracting model entirely but to create a new hybrid contracting model. This new model can leverage both the scalable, subscription-based framework from SaaS contracts and BPO-style performance and governance commitments. In doing so, it can allow both companies and providers achieve new levels of value in contracts for services delivered using agentic AI.

AUTHORS

PARTNER

ROHITH P. GEORGE

NORTHERN CALIFORNIA (PALO ALTO & SAN FRANCISCO)
+1 650 331 2014

RGEORGE@MAYERBROWN.COM

PARTNER

JOE PENNELL

CHICAGO +1 312 701 8354

JPENNELL@MAYERBROWN.COM

PARTNER

BRAD L. PETERSON

CHICAGO +1 312 701 8568

BPETERSON@MAYERBROWN.COM

PARTNER

OLIVER YAROS

LONDON +44 20 3130 3698

OYAROS@MAYERBROWN.COM

Mayer Brown is a global legal services provider comprising associated legal practices that are separate entities, including Mayer Brown LLP (Illinois, USA), Mayer Brown International LLP (England & Wales),

Mayer Brown Hong Kong LLP (a Hong Kong limited liability partnership) and Tauli & Chequer Advogados (a Brazilian law partnership) (collectively, the “Mayer Brown Practices”). The Mayer Brown Practices are established in various jurisdictions and may be a legal person or a partnership. PK Wong LLC (“PKW”) is the constituent Singapore law practice of our licensed joint law venture in Singapore, Mayer Brown PK Wong Pte. Ltd. More information about the individual Mayer Brown Practices and PKW can be found in the [Legal Notices](#) section of our website.

“Mayer Brown” and the Mayer Brown logo are the trademarks of Mayer Brown.

Attorney Advertising. Prior results do not guarantee a similar outcome.

APRIL 06, 2026

AI AS STAR WITNESS: HOW A BUYER'S AI CONVERSATIONS SANK ITS EARNOUT AVOIDANCE STRATEGY

AUTHORS: JONATHAN A. DHANAWADE, FRANK J. FAVIA JR., ANDREW J. STANGER

This article was republished by the [Harvard Law School Forum on Corporate Governance](#).

On March 16, 2026, the Delaware Court of Chancery issued a significant post-trial opinion in *Fortis Advisors, LLC v. Krafton, Inc.*¹ The case arose from Krafton, Inc.'s (the "Buyer") acquisition of Unknown Worlds Entertainment (the "Target"), and the Buyer's subsequent attempt to engineer its way out of a \$250 million earnout obligation. The court found that the Buyer breached the parties' Equity Purchase Agreement (the "EPA") by terminating the Target's three key executives without (contractually defined) "Cause" and by seizing operational control of the studio. In a sweeping remedial order, the court reinstated the Target's CEO to his position with full operational authority, enjoined the Buyer from using the studio's own board to circumvent that authority, restored the Target's access to its publishing platform, and equitably extended the earnout measurement period by 258 days. Perhaps most remarkably, the court quoted at length from responses generated by a popular AI platform in response to queries made by the Buyer's CEO about how to take control of the Target's operations.

The decision carries important implications for acquirers, sellers, and their counsel in M&A transactions, particularly with respect to the evidentiary risks of using AI tools for legal advice and strategy, for-cause termination rights, the negotiation and enforcement of earnout provisions, and the significance of specific performance provisions in M&A agreements.

BACKGROUND

In October 2021, the Buyer, a South Korean gaming conglomerate, acquired the Target, which created the *Subnautica* gaming franchise. *Subnautica* is an underwater survival adventure game that became a cultural and commercial phenomenon, selling over 17.5 million copies and exceeding \$300 million in gross revenue. The Buyer purchased the Target for \$500 million upfront, plus up to \$250 million in contingent earnout payments tied to post-closing revenue. To close the deal, the Buyer made a critical promise: the Target's two co-founders and its CEO (collectively, the "Key Employees"), would retain operational control of the studio in all material respects during the earnout period. This included authority over the product roadmap, launch timing, planning, partnering, budgeting, and employee matters. That control would endure so long as any one of the three remained employed.

In exchange, the Key Employees agreed to run the Target in the ordinary course of business, consistent with its historical customs and practices, and accepted that the Buyer's consent would be required for certain significant corporate actions. The EPA also priced in the possibility that a Key Employee might

leave during the earnout period: each voluntary departure would reduce the revenue credited to the earnout calculation by \$1 million per departing employee. The EPA included a narrow definition of “Cause” for termination: it was limited to felonies, intentional fraud or dishonesty, gross misconduct, and wrongful disclosure of trade secrets or confidential information, each requiring a high degree of culpability.

After the acquisition closed, the founders gradually transitioned away from day-to-day game development in ways that the Buyer knew about and, in many instances, actively supported. Both founders voluntarily reduced their salaries from nearly \$400,000 to \$100,000 to reflect their scaled-back roles; those changes were entered directly into the Buyer’s centralized HR system and reviewed by the Buyer’s finance team. The Target’s CEO, by contrast, continued to shoulder the full weight of studio operations, serving as both CEO and the Target’s primary liaison to an increasingly difficult parent company.

By mid-2025, *Subnautica 2* was on track for an early access launch in August, and the Buyer’s own internal projections indicated that a successful launch would generate revenue well above the earnout threshold, producing a projected payout of approximately \$191.8 million in the base case and \$242.2 million in the best case. The Buyer’s CEO felt that the Buyer had agreed to a “bad deal” and that paying the earnout would make him look like a “pushover.” After the Buyer’s own legal team warned him that a for-cause termination would not eliminate the earnout obligation and would expose the Buyer to legal and reputational risk, the CEO turned to an AI chatbot for help. When the chatbot told him the earnout would be “difficult to cancel,” the CEO complained to colleagues that the EPA was “a contract under which we can only be dragged around.” At the chatbot’s suggestion, the CEO formed an internal task force dubbed “Project X,” whose mandate was to either renegotiate the earnout or execute a full “takeover” of the Target. The AI chatbot furnished the CEO with a detailed “Response Strategy to a No-Deal’ Scenario,” which included a “pressure and leverage package,” strategic talking points for negotiating with the Key Employees, instructions to lock down the Target’s publishing platform and control publishing rights, directions to prepare systematic materials for legal defense, and a “two handed strategy” combining hardball legal and financial pressure with softer retention incentives. The CEO admitted at trial that he had deleted certain of the relevant logs from the AI platform.

The Buyer followed most of the chatbot’s recommendations. It locked the Target out of the publishing platform, effectively severing the studio’s ability to release *Subnautica 2*. It posted a message on the websites of the Target and *Subnautica*, which was drafted overnight and without the studio’s involvement, falsely claiming that the founders were considering an invitation from the Buyer to reengage with the project. When negotiations over the earnout stalled in late June 2025, the Buyer determined to execute its takeover strategy. On July 1, 2025, the Buyer sent termination letters to the Key Employees, citing a single ground for dismissal: their intention to proceed with a premature release of *Subnautica 2*. The Buyer simultaneously removed the Key Employees from the Target’s board, replacing them with Buyer representatives, and installed a part-time replacement CEO for the Target who had never played a *Subnautica* game and had no experience overseeing early access title development.

THE COURT’S ANALYSIS

THE TERMINATIONS WERE NOT FOR “CAUSE”

Vice Chancellor Will of the Delaware Chancery Court found that none of the Buyer’s proffered justifications for the terminations satisfied the EPA’s definition of “Cause.” Under the EPA’s narrowly negotiated definition of Cause, the only provision at issue was whether the Key Employees had committed an “intentional act of . . . dishonesty.” The court interpreted that phrase strictly, holding that

intentional dishonesty requires a conscious objective to deceive, not merely a deliberate act that might be characterized as inaccurate or imprudent.

At trial, the Buyer had abandoned its original stated basis for termination (that the Key Employees intended to prematurely release *Subnautica 2*) and advanced two new theories. The Buyer's first new theory at trial was that the founders had secretly entered "semi-retirement" and concealed their diminished roles from the Buyer. The court rejected this outright. The record showed that the Buyer had known about, and in many cases actively endorsed, the founders' evolving roles for over a year before the terminations. The Buyer's own internal communications reflected its understanding that one of the founders was not working on *Subnautica 2* directly and that his salary had been reduced. The evidence showed no deception, no coverup, and no intent to mislead.

The Buyer's second new theory was that the Key Employees' mass downloads of company data constituted terminable offenses for Cause. The court acknowledged that the downloads were "wrongful," but found they were defensive measures taken in good faith to protect the studio's work product in the face of the Buyer's escalating takeover, not acts of intentional dishonesty. Critically, the Key Employees kept the data confidential and returned it promptly after their terminations. As the court put it: "These are not the actions of thieves."

The court also rejected these new theories by applying the "mend-the-hold" doctrine, which barred the Buyer from relying on justifications it had known about before the terminations but chose not to cite in its termination notices, and the "after-acquired evidence" doctrine, which would have required the Buyer to prove that the later-discovered conduct was so severe that it would have independently resulted in termination if known at the time of discharge.

Because neither of the Buyer's new theories established the conscious objective to deceive required by the EPA, the court held that the Buyer lacked Cause to terminate the Key Employees.

THE BUYER WRONGFULLY SEIZED OPERATIONAL CONTROL

The court separately held that the Buyer's seizure of operational control was independently wrongful. The EPA conditioned the Key Employees' right to operational control on three requirements: (a) operating the Target in the ordinary course of business; (b) not taking certain enumerated actions without the Buyer's prior consent; and (c) exercising operational control in good faith and in material compliance with applicable law. The Buyer argued that the Key Employees had breached all three. The court disagreed with respect to each.

On the **ordinary course covenant**, the Buyer argued that the Target could not operate "consistent with past custom and practice" without its founders personally leading game development—i.e., that the founders were the essence of the studio's "magical pair" creative model. The court rejected this framing. It held that an ordinary course covenant is assessed at the company level, not the individual employee level. Before the acquisition, the founders had already demonstrated a practice of delegating primary development duties to other team members, most notably on *Subnautica: Below Zero*, which the founders did not lead. After the acquisition, the studio continued to design, develop, and prepare to publish games using its established community-focused, early access development model. The Buyer could not manufacture a covenant breach by pointing to changes in individual executives' job duties when the company's fundamental business operations remained unchanged.

On the **consent-required actions**, the Buyer advanced two theories. First, it argued that the founders' voluntary role and salary reductions triggered a "constructive termination" under their employment agreements, which required the Buyer's advance consent under the EPA. The court found this argument

commercially unreasonable: the founders could not create a constructive termination claim against themselves by voluntarily reducing their own salaries or modifying their own duties. Second, the Buyer argued that elevating game developers to project-level leadership constituted the hiring of “executive-level” employees requiring the Buyer’s consent. The court disagreed and found that these were routine personnel decisions squarely within the Key Employees’ retained operational authority, not the hiring of corporate executives.

On the **good faith obligation**, the Buyer argued that the Target’s CEO had encouraged the founders to “stay on the books” in reduced roles specifically to preserve the earnout and that this constituted bad faith. The court was not persuaded. The EPA did not obligate the founders to remain in their exact pre-acquisition roles indefinitely; it merely provided that operational control would persist so long as any one Key Employee remained employed. If both founders had departed, the only consequence would have been a \$1 million reduction per departing founder in the revenue counted toward the earnout threshold—a modest adjustment in the context of a potential \$250 million payout. In the court’s view, it was entirely reasonable for the Target’s CEO to encourage the founders to transition into reduced advisory roles rather than leave outright. Openly facilitating that transition was not bad faith.

Because the Buyer failed to prove that the Key Employees breached any of the conditions imposed upon their post-acquisition operational control, the court held that the Buyer breached the EPA when it removed their operational control.

THE REMEDY

Having found that the Buyer breached the EPA on two independent grounds (wrongful termination and wrongful seizure of operational control), the court awarded specific performance, finding that the Key Employees’ right to operational control was a crucial, bargained-for protection that money damages alone could not adequately replace. The EPA itself contained a specific performance clause, and the court held that the Buyer had offered no persuasive or case-specific reason to disregard it. The loss of creative control over a unique asset (the ability to guide the development and launch of *Subnautica 2*) constituted irreparable harm.

The court crafted a tailored remedy rather than a blanket restoration of the pre-termination status quo. Although the Target’s former CEO was reinstated with full operational authority, the founders were not restored to their pre-termination roles. The court found that restoring the Target’s former CEO alone was sufficient to vindicate the Key Employees’ contractual rights since they would still qualify for the earnout and the founders could still be engaged in an advisory capacity. The court declared that a previous resolution by the Buyer’s board, which had purported to prohibit the release of *Subnautica 2* absent board approval, was ineffective to the extent it infringed on the Target CEO’s operational authority. The Buyer was enjoined from using the Target’s board, or any other corporate mechanism, to circumvent the Target CEO’s decision-making authority over the game’s early access launch. The Buyer was also ordered to immediately restore the Target CEO’s access to the Target’s publishing platform.

To ensure the specific performance remedy was not rendered illusory by the passage of time, the court also equitably extended the earnout measurement period by 258 days, the exact duration of the Target CEO’s ouster measured from his July 1, 2025 termination to the date of the opinion. Whether the Buyer’s actions actually impaired the earnout, and what money damages (if any) are owed, remain reserved for a forthcoming second phase of the bifurcated litigation.

KEY TAKEAWAYS

AI communications may be discoverable and provide unique insight into a party’s motives. Perhaps

the most striking feature of the opinion is that the court quoted at length from the Buyer CEO's conversations with the AI chatbot and expressly relied on those exchanges to establish the pretextual nature of the Buyer's conduct. The chatbot's recommendations were treated as a window into the Buyer's strategic intent. The court also noted that the Buyer's CEO admitted at trial to deleting specific, relevant logs from the AI platform—a fact that may factor prominently in the second phase of the litigation, where money damages and potential earnout impairment are at issue. The opinion does not establish new legal standards governing AI use, but it does carry a clear warning: communications with AI platforms may not be privileged, may be subject to discovery, and may be used against the party that generated them. Deleting them may compound the problem. Acquirers using AI tools for post-closing strategy or any deal-related purpose should treat those communications with the same discipline they would apply to other non-privileged communications, including applying an appropriate retention policy.

The evidentiary and protected status (or lack thereof) of AI-generated communications is an area of law that is rapidly developing. A recent Mayer Brown Legal Update ([M&A Discovery in the AI Era: Generative AI Communications and Outputs May Become Litigation Ammunition](#)) examines in depth recent judicial opinions on the application of privilege to AI use. The landscape is also evolving legislatively: [a proposed New York bill introduced in March 2026](#) would bar AI chatbots from posing as lawyers and give deceived users the right to sue for damages. The *Krafton* case provides a vivid, real-world illustration of why practitioners and dealmakers should be paying close attention to these developments.

Be careful what you put in your termination notice because you will likely be bound by it. The *Krafton* case is a cautionary tale about what happens when, in the middle of litigation, an acquirer tries to substitute its stated basis for termination with new grounds. The court applied the mend-the-hold doctrine to bar the Buyer from relying on the founders' role changes as a justification for termination when the Buyer had known about those changes long before the terminations and chose not to cite them in its termination notices. For the data downloads, which the Buyer did not fully discover until after the terminations, the court applied the after-acquired evidence doctrine, which required the Buyer to prove that the downloads independently would have resulted in termination. The Buyer failed that test as well. More broadly, the court interpreted the EPA's "intentional act of dishonesty" standard as requiring a conscious objective to deceive, not merely deliberate conduct that could be characterized as improper. Acquirers should carefully consider whether their for-cause termination definitions are broad enough to cover the full range of conduct they may realistically encounter post-closing and should understand that the grounds cited in the termination notice are likely the grounds they will be stuck with at trial. Sellers should continue to push for narrow, specifically negotiated definitions of cause.

A breach during the earnout period can cost more than the breach itself. The *Krafton* decision demonstrates that Delaware courts may use their broad equitable powers to extend earnout measurement periods when a buyer's breach has deprived the sellers of the time and operational runway they were promised. In *Krafton*, the court extended the base earnout measurement period by 258 days—the exact duration of the improper ouster of the Target's CEO. The extension was calibrated to ensure that the specific performance remedy was not illusory: sellers need a genuine opportunity to achieve earnout targets, and a court-ordered reinstatement without time to generate the required revenue accomplished little. For acquirers contemplating post-closing interventions, the cost of a breach is not merely the price of lost litigation but is also the cost of additional time on the earnout clock.

Specific performance clauses carry real weight in Delaware. *Krafton* underscores that Delaware courts will enforce a specific performance clause unless the breaching party offers a "persuasive and case-specific" reason to override it. The court not only reinstated the Target's CEO but also issued broad

injunctive relief preventing the Buyer from using the Target's board of directors to circumvent the Target CEO's operational authority. The takeaway for sellers is that specific performance clauses are one of the most valuable protections in an earnout-heavy deal structure and should be negotiated robustly. The takeaway for acquirers is that the equitable remedies Delaware courts are willing to craft may extend well beyond simple reinstatement; they also encompass injunctions, earnout period extensions, and the nullification of board resolutions, leaving little room to maneuver post-breach.

¹ ___ A.3d. ___, C.A. No. 2025-0805-LWW (Del. Ch. March 16, 2026) (Will, V.C.).

AUTHORS

PARTNER

JONATHAN A. DHANAWADE

CHICAGO +1 312 701 7874

JDHANAWADE@MAYERBROWN.COM

PARTNER

FRANK J. FAVIA JR.

CHICAGO +1 312 701 8780

NEW YORK

FFAVIAJR@MAYERBROWN.COM

KNOWLEDGE COUNSEL

ANDREW J. STANGER

HOUSTON +1 713 238 2702

AJSTANGER@MAYERBROWN.COM

Mayer Brown is a global legal services provider comprising associated legal practices that are separate entities, including Mayer Brown LLP (Illinois, USA), Mayer Brown International LLP (England & Wales),

Mayer Brown Hong Kong LLP (a Hong Kong limited liability partnership) and Tauil & Chequer Advogados (a Brazilian law partnership) (collectively, the "Mayer Brown Practices"). The Mayer Brown Practices are established in various jurisdictions and may be a legal person or a partnership. PK Wong LLC ("PKW") is the constituent Singapore law practice of our licensed joint law venture in Singapore, Mayer Brown PK Wong Pte. Ltd. More information about the individual Mayer Brown Practices and PKW can be found in the Legal Notices section of our website.

"Mayer Brown" and the Mayer Brown logo are the trademarks of Mayer Brown.

Attorney Advertising. Prior results do not guarantee a similar outcome.

DECEMBER 19, 2025

ARTIFICIAL INTELLIGENCE PROVISIONS IN TECHNOLOGY CONTRACTING: KEEPING UP WITH THE EVOLVING REGULATORY LANDSCAPE

AUTHORS:

ANA HADNES BRUDER, JULIAN M. DIBBELL, GABRIELA KENNEDY, ARSEN KOURINIAN,
OLIVER YAROS

Source: Tech Talks Podcast

In this *Tech Talks* episode, Mayer Brown partners Ana Bruder, Julian Dibbell, Gabriela Kennedy, Arsen Kourinian, and Oliver Yaros put four AI regimes head to head: the European Union’s risk based Act, the United Kingdom’s light touch approach, Asia’s mixed models, and the United States’ state by state patchwork. Their contracting playbook: lock down roles and risk, and go beyond “comply with law” to require risk management, human oversight, transparency, and data use limits. Our hosts close out by looking ahead to 2026 to see how these regimes diverge further.

SHOW NOTES:

00:04 Introductions to Artificial Intelligence Provisions in Technology Contracting
02:52 EU AI Act: Roles, Risk Tiers, and Contract Implications
04:05 UK’s Pro Innovation, Light Touch Framework
07:50 Asia Snapshot: China’s GenAI Rules vs. Soft Touch Regimes
10:58 U.S. Patchwork: State Laws, High Risk Use Cases, and Contracts
14:41 Negotiating AI Contracts Amid a Shifting Global Regulatory Landscape
18:44 Practical Best Practices for AI Contracting
24:50 Beyond “comply with law”: Standards, Data Use, and Future Proofing
27:25 2026 Looking Ahead: Global Regulatory Trajectories

To learn more about this topic, see our recent article [Artificial Intelligence Provisions in Technology Contracting: Keeping Up with the Evolving Regulatory Landscape](#).

SUBSCRIBE

Listen to more episodes of [Tech Talks](#).

Tech Talks is also available on the following podcast services for your subscription convenience.

AUTHORS

PARTNER

JULIAN M. DIBBELL

CHICAGO +1 312 701 8237

JDIBBELL@MAYERBROWN.COM

PARTNER

ANA HADNES BRUDER

FRANKFURT +49 69 7941 1778

ABRUDER@MAYERBROWN.COM

PARTNER

GABRIELA KENNEDY

HONG KONG +852 2977 1790

GABRIELA.KENNEDY@MAYERBROWN.COM

PARTNER

ARSEN KOURINIAN

LOS ANGELES +1 213 229 5141

AKOURINIAN@MAYERBROWN.COM

PARTNER

OLIVER YAROS

LONDON +44 20 3130 3698

OYAROS@MAYERBROWN.COM

Mayer Brown is a global legal services provider comprising associated legal practices that are separate entities, including Mayer Brown LLP (Illinois, USA), Mayer Brown International LLP (England & Wales), Mayer Brown Hong Kong LLP (a Hong Kong limited liability partnership) and Tauil & Chequer Advogados (a Brazilian law partnership) (collectively, the “Mayer Brown Practices”). The Mayer Brown Practices are established in various jurisdictions and may be a legal person or a partnership. PK Wong LLC (“PKW”) is the constituent Singapore law practice of our licensed joint law venture in Singapore, Mayer Brown PK Wong Pte. Ltd. More information about the individual Mayer Brown Practices and PKW can be found in the [Legal Notices](#) section of our website.

“Mayer Brown” and the Mayer Brown logo are the trademarks of Mayer Brown.

Attorney Advertising. Prior results do not guarantee a similar outcome.