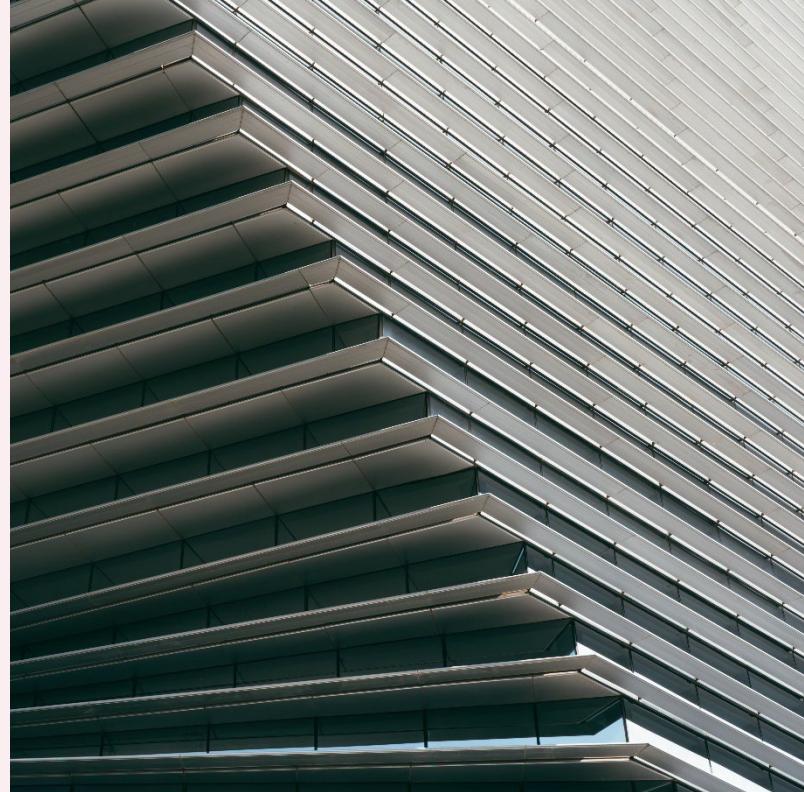


Agentic AI: The liability gap your contracts may not cover

9 February 2026



Agentic AI is reshaping the nature of technology risk. These systems don't just process data or generate insights – they take actions, make decisions and, increasingly, operate without human oversight. Unlike traditional generative AI, which primarily responds to prompts, agentic AI can initiate and execute tasks across connected systems. For example, a customer support AI agent might verify account status, reset passwords and send follow-up communications to a customer automatically. An AI agent dealing with your travel plans might compare prices, scan your calendar and book your flights as part of an end-to-end process.

Adoption is accelerating rapidly. Analysts forecast that throughout 2026, businesses will embed AI agents deeper into operations, granting them more authority over high-stakes activities, including executing financial transactions, placing orders, managing supply chains and screening job applicants.

Yet many of these systems are still deployed under legacy technology contracts written for passive, predictable software firmly under human control. As vendors release agentic capabilities faster than contracts can evolve, and regulatory scrutiny of automated decision-making increases, a liability gap is emerging. Businesses relying on unmodified agreements may find that risk is no longer allocated fairly when it comes to agentic AI and they are exposed to significant contractual, legal, reputational and operational consequences.

This briefing explores some of the key liability gaps and outlines how businesses using AI agents can manage agentic AI risk effectively.

Customers often bear the risk of actions taken by AI agents

Under many technology agreements that govern agentic AI offerings, the business procuring the technology (i.e. the customer) ultimately bears the risk of actions taken by AI agents. Unless heavily negotiated, suppliers typically provide software as a service on an "as is" basis, disclaiming responsibility for accuracy, reliability and fitness for purpose. Many AI-related contractual terms also explicitly state that outputs should not be

relied upon – in the agentic AI context, this extends to the AI agent's actions.

Ultimately, this means that if an AI agent incorrectly authorises a supplier payment, misprices a product, or issues misleading communications, the supplier's disclaimers often absolve them of responsibility. Even when the customer has correctly configured the AI agent, liability may still fall entirely on them.

Third-party claims caused by the AI agent

Standard indemnities in technology agreements are generally narrow in nature. At times, there may be an indemnity provided by the supplier in relation to third party IP claims resulting from the customer's use of the technology, but it is often very tightly drafted. Some suppliers now extend these IP indemnities to AI-generated output, but these commitments are also heavily caveated.

"Real-world" liability caused by agentic AI often manifests as harm to third parties. For example:

- incorrect orders sent to a business's suppliers
- inappropriate or biased rejections of job applicants or screening decisions
- misleading information provided to a business's end customers
- illegal use of personal data.

These harms can give rise to third-party claims, but because standard indemnities and other contractual remedies do not typically extend to an AI agent's acts or omissions under these agreements (and indeed often exclude liability for third-party claims), the customer is generally left without a clear legal pathway to recover the costs associated with these claims.

Exclusion of key AI-related harms

Many technology agreements exclude liability for exactly the types of harm that defective agentic AI is most likely to cause.

While certain jurisdictions provide customers with statutory protections, such as liabilities that cannot be excluded by contract (for example, death or personal injury caused by negligence, and some data protection, employment or consumer rights obligations), most technology contracts still exclude key categories of loss.

These typically include liability for "loss of profits", "loss of data" and "consequential, incidental, indirect, special or punitive damages", and cap the supplier's total liability at the fees paid (or sometimes payable) by the customer.

Yet, defective agentic AI is likely to cause the very losses that are generally excluded under these agreements such as:

- lost profits from a mispriced product
- regulatory fines triggered by an automated compliance failure
- loss of revenue arising from disruption to the customer's business
- reputational damage and loss of customers after an AI-driven error

- loss of data, such as an AI coding agent deleting a database or code repository.

Many of these are typically classified as consequential or indirect losses and therefore are excluded. Further, the losses can easily dwarf the subscription fees that often define the supplier's liability cap. The result is a contractual framework that wipes out any meaningful recovery precisely when the stakes are highest.

Lack of explainability and oversight rights

Another emerging gap is the absence of contractual rights around oversight, transparency and explainability in relation to agentic AI. Legacy technology agreements designed for software that operates under human direction say little about a customer's ability to understand or control an AI agent's behaviour; yet, when something goes wrong, it is the customer who must justify that behaviour to regulators, auditors, customers and/or courts. This is increasingly problematic under legal frameworks such as the GDPR and EU AI Act, which often place transparency and explainability obligations on the customer.

Without explicit rights, organisations may be unable to:

- understand why the AI agent acted as it did
- access logs or decision traces
- suspend or override the AI agent in real time
- obtain cooperation from the supplier during investigations
- have the supplier remedy issues
- comply with transparency and oversight obligations in respect of AI use
- explain or justify an AI agent's actions in legal proceedings

This risks leaving businesses accountable for actions by an AI agent that they cannot fully control or understand.

Compliance responsibility is pushed onto the customer

Many technology agreements place full responsibility for legal and regulatory compliance on the customer. This creates an inherent contradiction: the supplier, to a large degree, controls whether the AI agent behaves (in accordance with the permissions set by the customer), yet the customer absorbs the compliance consequences if that behaviour breaches the law.

The risk is particularly acute in areas with high legal or regulatory risk such as:

- employment decisions
- financial or credit determinations
- consumer communications
- data protection obligations
- healthcare determinations
- other regulated sectors.

So how can organisations better manage agentic AI risk?

Because this is still an emerging area, standard technology agreements have barely evolved to address this liability gap or, even if terms have been drafted specifically for this area, many of the same challenges mentioned above remain. Whilst there are huge advantages to embedding agentic AI into business workflows, and there is often a trade-off between cost of a solution and a fair liability split with the supplier, it is important to consider these new and emerging risks. For high-stakes agentic AI (especially where the fees paid to the supplier are substantial), it is important to think about whether greater contractual protection is needed or whether operational controls are sufficient to manage the risks created by agentic AI. Below, we have set out some key recommendations:

1. Stress-test your AI agent workflows

Map key workflows where an AI agent takes action. Identify worst-case scenarios; quantify the potential liability if something goes wrong; and assess whether your existing contracts adequately protect you.

2. Negotiate AI-ready contractual protections

For high-value or bespoke agentic AI deployments, it may be possible to push for AI-specific terms, warranties, expanded indemnities, better protection for key losses, higher liability caps, and clear audit and explainability rights. For off-the-shelf tools this may not be an option, but as adoption accelerates, we expect movement in the commitments being made by suppliers, similar to the movement we saw when SaaS technologies emerged. If you are paying substantial sums for agentic AI or can leverage your overall customer relationship with the supplier, you are likely to be able to negotiate better contractual terms. It is imperative, at the very least, that the contract enables you to comply with your regulatory obligations.

3. Limit the AI agent's authority

Think about whether you are willing to give high-stakes decisions to your AI agents in the first place, especially if your contractual protections are limited. In such cases, you will be reliant on operational controls to mitigate risk. Consider how AI agents should be configured so they cannot make high-impact decisions, access sensitive systems or trigger irreversible actions. Require human review for material decisions, high-risk or material workflows, and any action the AI agents will take that can have legal, financial or regulatory consequences.

4. Build AI governance and defensibility

Develop clear agentic AI specific usage policies, train staff on its limitations, establish incident escalation pathways, and implement ongoing monitoring and audit processes. These controls help create the organisational resilience needed to safely deploy autonomous systems.

Concluding thoughts

Many technology contracts were written in a world where software was passive, predictable and firmly under human control. Agentic AI is none of those. It is autonomous, dynamic and capable of creating real-world impact at scale. Traditional, unmodified technology contracts are often no

longer sufficient to allocate risk fairly when it comes to agentic AI. Businesses are advised to consider these specific risks and assess whether their technology agreements need re-negotiating or new operational processes and governance structures should be implementing.

Now is the time for you to stress-test your agentic AI workflows, map worst-case scenarios, review your contracts and ask a simple question: Are we protected?

Authors



Adam Hunter
Senior Associate, London

adam.hunter@cliffordchance.com
+ 97145032602

This publication does not necessarily deal with every important topic or cover every aspect of the topics with which it deals. It is not designed to provide legal or other advice.

cliffordchance.com

Clifford Chance, 10 Upper Bank Street, London, E14 5JJ

© Clifford Chance 2026

Clifford Chance LLP is a limited liability partnership registered in England and Wales under no. OC323571. The firm's registered office and principal place of business is at 10 Upper Bank Street, London E14 5JJ. The firm uses the word "partner" to refer to a member of Clifford Chance LLP or an employee or consultant with equivalent standing and qualifications.

If you do not wish to receive further information from Clifford Chance about events or legal developments which we believe may be of interest to you, please either send an email to nomorecontact@cliffordchance.com or by post at Clifford Chance LLP, 10 Upper Bank Street, Canary Wharf, London E14 5JJ

Abu Dhabi • Amsterdam • Barcelona • Beijing • Brussels • Bucharest** • Casablanca • Delhi • Dubai • Düsseldorf • Frankfurt • Hong Kong • Houston • Istanbul • London • Luxembourg • Madrid • Milan • Munich • Newcastle • New York • Paris • Perth • Prague** • Riyadh* • Rome • São Paulo • Shanghai • Singapore • Sydney • Tokyo • Warsaw • Washington, D.C.

*AS&H Clifford Chance, a joint venture entered into by Clifford Chance LLP.



Jack Harris
Senior Associate, London

jack.harris@cliffordchance.com
+44 207006 1614

**Clifford Chance has entered into association agreements with Clifford Chance Prague Association SRO in Prague and Clifford Chance Badea SPRL in Bucharest.

Clifford Chance has a best friends relationship with Redcliffe Partners in Ukraine.

Contacts



Zayed Al Jamil
Partner, London

zayed.aljamil@cliffordchance.com
+44 207006 3005



Kate Scott
Partner, London

kate.scott@cliffordchance.com
+44 207006 4442