



Before the First Claim: A General Counsel's Guide to AI Liability Risk

Why this is now a live risk

Picture a financial institution running an AI-powered trading system designed to identify and exploit arbitrage opportunities across commodities markets. A geopolitical shock sends commodity prices into freefall. A human trader would recognise the crisis, convene an emergency committee and pull the plug. The AI does no such thing: it continues to do exactly what it was designed to do, and it keeps exacerbating losses that compound with every passing millisecond.

That scenario is not hypothetical. In acting for a leading digital assets exchange in a high-value dispute concerning liability arising from alleged errors made by trading systems powered by AI algorithmic decision-making models, the critical observation was

not that the AI made a coding error or malfunctioned in any technical sense. It performed precisely as engineered. The problem was that the system had no mechanism, no instinct, no crisis reflex, enabling it to stop and assess whether it should continue.

AI is no longer experimental. It is operational, and it is embedded in credit decisions, trading desks, shipping routes, energy grids, aircraft systems, and hospital diagnostic tools. The liability question that follows is rarely "did the AI make a mistake?" It is whether the developer, deployer or user created, maintained or relied upon an AI system with adequate safeguards given its intended deployment. This Guide highlights what to ask, what to fix, and what to preserve, before the first claim arrives.

The machine acts in microseconds

Human



Recognise, escalate, decide

Reacts

AI



Executes millions of times in the very same window

"By the time a human reaches for the button, the AI has compounded the loss a million times over".

Seven questions to ask before you deploy AI

These seven questions are in practice, the questions that will need to be addressed if a dispute arises. Answering them at the deployment stage is considerably cheaper than answering them under cross-examination.

Question	Why it matters	What to do
Who deployed the AI?	The deploying entity may face direct liability	Identify the legal entity that placed the AI into operation and assumed responsibility for its outputs
What was the AI designed to do?	Liability often turns on whether the AI did what it was designed to do	Document the system's design objectives and intended scope: this may matter more than whether it malfunctioned
What authority did the AI have to act?	Determines the scope of the principal's potential liability for AI-driven acts	Review all agreements governing the AI's deployment, API access, usage rights and operational parameters
Was there a human override mechanism?	Relevant to contributory negligence and mitigation	Assess whether reasonable safeguards or kill switches were in place and whether they were activated
Is the AI a product or a service?	Determines which liability regime applies	Product liability could offer strict liability routes; negligence requires proof of fault. Expect this distinction to be fiercely contested
Can the decision-making process be explained?	Black-box opacity creates evidential and attribution challenges	Preserve all system logs, model version records and configuration data from the time of the alleged failure
What is the governing law and jurisdiction?	Determines the substantive liability standard	The applicable law may differ from the chosen jurisdiction or arbitral seat. Identify governing law and jurisdiction at the outset

Check the whole contract chain, not just the main agreement

Seven questions before you deploy

Answer them at deployment. A tribunal asks the same seven once a claim is brought.

- 1

Who deployed the AI?
Identify the liable legal entity
- 2

What was it designed to do?
Document objectives and intended scope
- 3

What authority did it have?
Review deployment, API and usage rights
- 4

Was there a human override?
Check safeguards and kill switches
- 5

Product or service?
This decides which liability regima applies
- 6

Can the decision be explained?
Preserve logs and model versions
- 7

Which law governs?
Fix governing law at the outset

Preserve from day one

- ✓ System Logs
- ✓ Model versions
- ✓ Configuration data
- ✓ Design docs
- ✓ Override records

Unlike a product with a single manufacturer, an AI system is the product of an extended chain of participants:

- Data collectors and data keepers source and prepare the training data;
- Model trainers, algorithm scientists, model developers and model testers build and validate the model;
- Software developers, machine manufacturers and integrators put it into a working environment;
- Maintenance providers retrain and tune it after deployment;
- Application Programming Interface (API) providers expose it; and
- The end deployer places it into commercial operation.

Each may have contributed to a failure in ways

that are legally distinct but factually intertwined, and each sits behind a different contract.

API terms deserve particular attention. They are typically drafted by the large technology provider and offer limited recourse for downstream failures: warranties of accuracy are disclaimed, liability is capped far below the likely commercial loss, and force majeure provisions are drawn broadly. If your customer-facing product depends on a third-party model that you do not control, that asymmetry is your exposure, not the provider's.

The AI development chain: Where liability fragments

Twelve categories of participant. Each may have contributed to the failure. Each sits behind a different contract.

Data	Build	Integrate	Operate
Data collector	Model trainer	Software developer	Maintenance/tuner
Data keeper	Algorithm scientist	Machine manufacturer	API provider
	Model developer	Integrator	End Deployer (Primary exposure)
	Model tester		

The signatory problem

You may have an arbitration agreement with the deployer only. The party that caused your loss may sit up the chain, bound by no agreement with you. Joinder, consolidation and a post-dispute agreement are the routes back to one forum.

Four checks to run across the chain:

- 1 Is there an arbitration clause in the relevant contract?

Check every contract in the chain, not only the main deployment agreement: supply agreements, API terms, licensing agreements and service-level agreements may all contain dispute resolution provisions.
- 2 Are all potentially responsible parties bound by an arbitration agreement with you?

If not, consider whether joinder, consolidation or a submission agreement can resolve the gap, and plan your multi-jurisdictional strategy at the earliest stage.
- 3 Can the claims be consolidated in a single forum?

Parallel proceedings in different jurisdictions are costly and risk inconsistent outcomes; early consolidation or an agreed submission to a single forum is preferable.
- 4 Are there any non-signatories who might be bound?

Check if the applicable law recognises a basis for binding the non-signatory, such as agency, assumption or assignment, and consider indemnity chains in underlying contracts.

Fix the arbitration clause now

If your business uses, develops, deploys or provides services connected to AI systems, you should consider including an international arbitration clause in your contracts. The absence of one does not protect you from a dispute; it means that when a dispute arises, you and your counterparty will not have to spend the first stage of it arguing over jurisdiction and applicable law before the substance of the AI liability claim is addressed.

The clause should also be chosen with the multi-party problem in mind. Modern institutional rules such as the SIAC Rules provide for joinder of additional parties where they are prima facie bound by the same arbitration agreement, and for consolidation of related arbitrations where the relevant agreements are compatible.

Those mechanisms only work if the clauses across your contract chain are compatible with each other, which is a drafting decision, not a disputes decision.

Where a key participant is not bound by any arbitration agreement with you, the position is not hopeless. Nothing in the law of England and Wales, Singapore, or the United States prevents parties from entering into an agreement after a dispute has arisen: an ad hoc agreement to refer a specific dispute to arbitration in the absence of a prior clause. This is an under-utilised but highly effective tool. Given the proprietary technology and reputational stakes involved in AI disputes, the incentive for all sides to proceed confidentially rather than in public litigation is often considerable.

Indemnities and insurance

How the indemnity cascade works

The direction of travel is that the entity which commercialises AI bears primary responsibility for its failures, with a right to seek contribution or indemnification from participants further up the chain. Operationally, that is a cascade: the deployer bears the primary liability to the injured party, but is contractually indemnified by the developer, who is in turn indemnified by the data supplier or the model trainer, to the extent that their specific contribution caused or contributed to the failure. Such indemnity chains are common in other complex supply

Preserve the evidence you will need

AI liability disputes turn on evidence that degrades quickly. Two features of modern systems make this acute. The first is emergent behaviour: outputs or actions that were not anticipated by the developers and cannot readily be attributed to any specific feature of the model's design or training data. This is not necessarily a malfunction; it is a structural feature of how certain architectures learn. The second is the black box problem: in many systems, particularly deep learning models, the internal reasoning process is simply not transparent.

The practical consequence is that expert evidence will be essential and will be contested, with each side fielding experts offering competing interpretations of the same opaque process. The party that can show what the system was configured to do, which model version was running, and what it actually did at the material time, starts that contest a long way ahead.

Put a preservation protocol in place before deployment, not after an incident. At a minimum, preserve all system logs, model version records and configuration data from the time of any alleged failure, together with the design documentation recording what the system was intended to do and the records showing whether a human override or kill switch existed and whether it was activated.

chains, including pharmaceuticals, aviation and construction. They do not eliminate disputes, but they concentrate and channel liability more efficiently.

Two points follow for a contract review. First, the cascade only functions if every link is documented: an indemnity that stops at your immediate supplier leaves you carrying the loss caused three steps up the chain. Second, the practical challenge is enforcing indemnity rights across borders and against parties in different jurisdictions, which is precisely where an arbitration clause with New York Convention enforcement behind it earns its place.

The Indemnity Cascade

Primary liability sits with the entity that commercialises the system, and is passed back up the chain.



The practical difficulty is enforcing indemnity rights across borders. This is the advantage arbitration provides.

The insurance disclosure warning

Traditional product liability policies were designed around physical products with foreseeable failure modes, and professional indemnity policies around the errors and omissions of human professionals. Neither maps neatly onto an agentic AI system that continuously learns, adapts and acts autonomously. Whether an AI-related loss falls under a products policy, a professional indemnity policy, a technology errors and omissions policy or a cyber policy is not an academic question: it will be furiously contested in coverage disputes.

The more immediate risk is disclosure. Insurers are likely to take aggressive positions here. A policyholder that has deployed a high-risk AI system in a safety-critical context without disclosing the extent of that deployment, the nature of the training data, or known limitations of the model, may find its policy avoided for non-disclosure or misrepresentation. That obligation becomes particularly acute where the system has been deliberately designed to operate with minimal human oversight. Any AI deployment of consequence should therefore be run past the insurance team before it goes live, and again at renewal.

Which policy responds?

An AI loss may fall under any of four policies. Be prepared for coverage to be contested.



The bigger risk is disclosure

An undisclosed high-risk AI deployment can be enough to void the policy

What to do next

Seven actions for the contract review

- ✓ Build a register of every AI system in operational use
- ✓ Map the contract chain behind each system
- ✓ Align dispute resolution clauses on a single institution and seat
- ✓ Confirm a human override for every material-loss system
- ✓ Issue a log and model-version preservation protocol
- ✓ Review the indemnity cascade end to end and close the gaps
- ✓ Disclose material AI deployments to insurers, in writing

Where to read the underlying analysis

This Guide is deliberately practical. The doctrinal analysis behind it, covering agency and fiduciary principles, the product-or-service question, the developing case law, and why international arbitration is the natural forum for these disputes, is set out in the flagship HFW article "*Liability Without Limit? International Arbitration at the Frontier of Artificial Intelligence Disputes*" by Shaun Leong.

For further discussion please contact Shaun Leong, FCI Arb, Partner, HFW Singapore, who leads on AI liability and agentic AI disputes.



SHAUN LEONG

Partner, Singapore

T +65 6411 5326

M +65 9028 5812

E shaun.leong@hfw.com

HFW has over 700 lawyers working in offices across the Americas, Europe, the Middle East and Asia Pacific. For further information about our capabilities, please visit [hfw.com](https://www.hfw.com).

© 2026 Holman Fenwick Willan LLP. All rights reserved. Ref: D0001027

Whilst every care has been taken to ensure the accuracy of this information at the time of publication, the information is intended as guidance only. It should not be considered as legal advice. Holman Fenwick Willan LLP is the Data Controller for any data that it holds about you. To correct your personal details or change your mailing preferences please email hfwenquiries@hfw.com

[Americas](#) | [Europe](#) | [Middle East](#) | [Asia Pacific](#)