




WHAT MATTERS NOW

The Wild West of AI Risk

Candid perspectives from across the market on a risk the industry is still defining.

About

is redefining how the insurance industry understands cyber risk in the age of AI.

Built specifically for insurance, transforms internet-scale cyber risk data into clear, actionable intelligence that supports smarter underwriting, portfolio management and exposure oversight. By combining internet-scale asset discovery, cyber risk assessment and continuous monitoring, helps insurers understand not just where cyber risk exists, but how it is evolving over time.

With visibility into any organisation with an online presence, including 100% coverage of the SMB sector, enables insurers to assess cyber risk at a scale and depth that traditional approaches often cannot achieve.

As AI reshapes both cyber threats and insurance workflows, provides the trusted intelligence layer insurers need to understand, quantify and manage cyber risk with confidence, while helping the market prepare for emerging AI-related exposures.

Headquartered in London, partners with leading insurers, brokers and reinsurers worldwide to make cyber risk understandable, actionable and insurable.





Table of contents

■ Introduction	04
■ What we mean by AI risk	05
■ Should AI be its own line?	06
■ Where the market agrees	08
■ Where the market splits	10
■ Shadow IT, shadow AI	12
■ The regulatory picture	13
■ However you cut it, you need visibility	14
■ The view	15

Introduction

AI has entrenched itself in business at remarkable speed. It is already drafting communications, screening candidates, and advising customers that used to sit with a person, and every one of those uses carries exposure. In each, AI is not only doing the work but becoming something that can make the wrong call or be turned into a way in – screening out a candidate it should never have rejected, giving a customer wrong advice, or being fed manipulated data until it becomes the weak link a claim eventually traces back to. The market knows this. What it does not yet have is an agreed way to name, price or cover the risk.

Beneath it all sits one unresolved question: is AI a risk in its own right, something that needs its own line of cover, or is it just a new type of risks the market already writes? The answer shapes everything downstream, from which policy a claim lands in to whether anyone is pricing for it at all. Yet the market has not settled it.

We went to the market to find out where thinking actually stands, speaking with carriers, brokers and specialty managing general agents (MGAs) writing AI cover from scratch. What came back was a set of well-argued, often conflicting positions, and a few points almost everyone shared.


This report sets out both: where the market agrees and where it splits – and what that means for anyone underwriting AI today.

What we mean by AI risk

First things first – let’s explore what people actually mean when they say ‘AI risk’, because even that is contested. There is no agreed definition, but the conversations point to a workable one. AI risk is the exposure a business takes on when it hands tasks, decisions or judgement to an AI system, whether or not anyone has decided to take that exposure on.

What makes it its own kind of problem is where the harm comes from. Some of it looks like cyber: a model is attacked, manipulated or leaks data. But much of it arrives when the system does exactly what it was built to do and gets the judgement wrong, which looks far more like professional advice gone bad than a breach. A recruitment tool screens out a candidate it should not have. A customer-facing chatbot promises something the business never intended. A model trained on the wrong material reproduces someone else’s work.

The harms that came up again and again were the same across conversations:



The common thread is that the harm rarely stays in one place. A single AI failure can trip several policies at once, from cyber to employment practices to general liability to directors and officers.

That breadth is exactly why the market cannot agree on where AI risk belongs, and it is the question we turn to next.

Should AI be its own line?

Ask the market whether AI deserves its own line of cover, and opinions split cleanly down the middle.

That question underlies pricing, wordings and exclusions alike, so it is worth addressing directly before going further. Neither side is clearly right; the same facts support both conclusions.

The case for a separate line

One camp sees something genuinely new, a risk that does not sit comfortably in any existing wording and will eventually need its own. The argument runs that AI exposure behaves unlike the perils current policies were built for. The loss often comes from the system working as intended and making a judgement that goes the wrong way, which is closer to professional indemnity than to cyber.

On this view, AI may follow the path cyber took: a risk first absorbed into existing policies, then gradually pulled out into a defined product of its own once the market understood it well enough to name and price it. A separate line would force the disclosure, the wordings and the pricing discipline that scattered coverage never quite delivers. The specialists already building standalone cover make the same case from the front line, arguing the existing forms were never designed for this exposure in the first place.

“

We view AI agents as effectively a new form of professional indemnity. A lot of the downside is when the system functions exactly as it was intended to. It just made a judgement call that went the wrong way.

Charles Higgins

Co-Founder & CTO, Redberry Labs

“

Existing policies were never intended to cover generative AI, and ultimately it is going to need a new category, with specialty underwriters to understand it, price it and bring it to market.

George Lewin-Smith

CEO & Co-Founder, Testudo

The case for keeping it where it is

The other camp sees no new peril, just more software on the network, already picked up by existing wordings. On this read, an AI tool is one more asset doing one more job, and the exposures it creates – a data leak, a faulty output, a liability to a third party – are ones the market already covers under cyber, technology E&O, or general liability. Carving out a separate line risks inventing a category the losses do not actually need, and creating gaps and overlaps between the new product and the old ones.

A harder-edged version of the same instinct is to exclude rather than define. Much of the market has moved first by writing the risk away, and in the US that is now built into the standard forms, with general liability exclusions for AI that carriers can switch on at will.

“

I think the only thing the industry is keeping up on is excluding AI from areas it doesn't want exposure to. In the US, you've seen a lot of general liability and ISO forms introducing AI exclusions – very similar to what they did when cyber became a thing, taking general liability off the table. Where we haven't quite settled is in the E&O and cyber space: how it fits, what would actually be triggered.

MGA

Why the question stays open

Both arguments hold up, and they lead to very different books. The reason the market cannot close the question is that the disagreement is not really about product design. It runs deeper, into what kind of risk AI even is, what to do about it, and whether it can be assessed at all before the claims arrive. Those are the fault lines that decide the answer, and they are where we go next, after the common ground that sits beneath all of it.

Where the market agrees

For all the disagreement, a layer of common ground runs underneath it, and it is worth starting there. A few points came up again and again, whoever we asked.

The real danger is the AI nobody manages

The exposure people fear most is the AI nobody is watching. It rarely comes through the front door. A marketing team adopts a tool off the shelf, an employee pastes sensitive data into a chatbot, a workflow quietly hands a decision to a model, and none of it is declared.

“

The biggest risk is when AI becomes unmanaged and people don't understand how it's being deployed or used within their systems, and how their employees are inputting data and using it.

Rachel Rossini

Head of Middle Market –
Cyber and Technology E&O, AXA XL

Most insurers have not even asked their insureds whether they use AI, so the exposure sits on the books unpriced. That blind spot, more than any single failure, is what the market keeps returning to.



The exposure runs well beyond cyber

There is broad agreement that AI risk does not stay in the cyber lane. The same AI failure can surface as an employment practices claim, a general liability claim and a D&O claim, each landing in a different policy.

This complicates the ‘keep it where it is’ case: even those who would not carve out a separate line accept the harm does not respect policy boundaries.

“

It’s a fundamentally different type of liability exposure. People who suggest AI is just cyber don’t have a good grip on how the threat percolates.

George Lewin-Smith
CEO & Co-Founder, Testudo

The claims are here, the data is not

Everyone agrees the claims have already started. Organisations are reporting losses linked to AI, but the long-term loss data insurers normally rely on simply does not exist yet.

“

There’s a lot of copyright infringement claims out there now, and data bias as well.

John Farley
Managing Director, Cyber Practice Group Leader, Gallagher

The sharper examples are doing the rounds too: an AI video-interview tool that allegedly rejected a deaf employee for a promotion and then advised her to ‘practice active listening,’ a chatbot prompt-injected into agreeing to sell a car for a dollar, or, in one report shared with us, an AI that opened an infected email link and spread malware across a company’s IT infrastructure. One specialist recounted a conversation with a broker who reported over \$150m of covered AI claims and more than \$1bn in uncovered ones sitting in his book. What most of the market does not yet have is a body of settled claims to price against. These cases take years to work through the courts, so the loss history underwriters would normally lean on is still some way off, and as one MGA put it, the data will not be there for another three or four years even as adoption keeps accelerating. A few specialists say they can already price it with confidence, by other means. Most are waiting for the losses to land.

Where the market splits

The consensus goes no further. Beyond it, opinion pulls apart on almost everything that matters – how to classify the risk, how to handle it, and whether it can be assessed at all before the claims arrive.

What kind of risk it is

This split decides the separate-line question, and it is the one the market is furthest from closing. The new-category and cyber-by-another-name camps set out earlier are not converging, and until one wins out, the other questions below stay open as well. It is worth noting how much rides on it: the same claim lands in a different policy, at a different price, depending only on which view the underwriter takes.

What to do about it

There is also little agreement on what to do about it. The exclusion-first approach described earlier is now the default position for much of the market, particularly in the US. Others argue the industry has to take the risk on rather than write it away.

“

We can't exclude AI from coverage. It's just the new way of operating. It would be like if it were 1995 and we were trying to exclude the internet from any sort of cyber claims.

Rachel Rossini

Head of Middle Market –
Cyber and Technology E&O, AXA XL



Whether you can assess it at all

The third split is whether you can assess it before the claims arrive. Insurance runs on disclosure, but a questionnaire only tells you what an insured chooses to share, and there is no equivalent yet of the outside-in tools the market leans on for cyber. For much of the market, assessing AI today means leaning on what the insured is willing to declare.

“

A lot of underwriters don't have the expertise to get into the black box of a platform. They ask what guardrails you have, whether there's a human in the loop monitoring the output. And often it isn't even your product, it's someone else's, so it comes down to what the contract says about liability when the thing goes sideways.

John Farley

Managing Director, Cyber Practice Group Leader, Gallagher

Some stay deliberately out of the technical detail anyway, on the view that only a small share of AI lawsuits actually turn on how the model performed. The suits that are landing tend to turn on how and where the system was built and used – copyright over training data, biometric-privacy claims, data exposure and discrimination – rather than on model accuracy. Others think the data is there if you go looking for it.

“

The nice thing about AI is that it is fundamentally a computer programme which is probabilistic, which means you can simulate it. I can't lock a legal professional in a room and ask them questions for years on end. I can do that to an agent.

Charles Higgins

Co-Founder & CTO, Redberry Labs


Either way, the gap is the same. The harm rarely stays in one place, so if you cannot see where AI is being used, you cannot see which of your lines it lands in.

Shadow IT, shadow AI

The exposure that unsettles underwriters most has a familiar shape. The market lived through silent cyber, where risk built up inside policies that were never written with it in mind, and the bill arrived later in claims nobody had priced. Carriers were burned by it once, and the wariness shows.

Silent AI is forming the same way, one undeclared tool at a time. Adoption is running ahead of disclosure. Businesses are wiring AI into hiring, customer service, claims handling and a hundred other workflows faster than their insurers can ask about it, and often faster than they are tracking it themselves. The exposure accumulates across a book without anyone deciding to take it on.

What makes it harder than shadow IT is concentration. The same handful of underlying models sit behind much of this activity, so a single flaw or failure can surface across many insureds at once. Risk that would once have been spread across dozens of separate systems now runs through a few shared ones. Caught early, it is manageable. Discovered after the fact, it starts to look like an accumulation event.



The regulatory picture

Regulation is starting to shape how AI risk gets defined and allocated, and it is moving at very different speeds on each side of the Atlantic.

The US picture

In the US, regulation is not yet doing the work that litigation is. There is no comprehensive federal AI statute, just a patchwork of state laws – and even that is unsettled. Colorado's first-in-the-nation AI law ([SB 24-205](#)) never took effect: after delays and a court-ordered stay, it was repealed in May 2026 and replaced by a lighter, disclosure-focused law ([SB 26-189](#)) that starts in January 2027. That does not mean AI exists outside the law: many claims are already being brought under existing frameworks, from discrimination and privacy to copyright and product liability. Narrower measures are live elsewhere – [Illinois's amended Human Rights Act](#) and [New York City's hiring bias-audit rule](#) among them – while the federal government is pushing back on state AI laws through [executive action](#). For now, the courts are shaping this risk faster than the legislatures, and US litigation culture is why most of the activity is concentrated there.

Across the Atlantic

Europe is moving towards formal rules, and they reach US and UK firms too.

From 2 August 2026, the [EU AI Act's](#) transparency obligations apply: businesses must disclose when people are dealing with an AI system, and label AI-generated content. The Act's tougher high-risk obligations have been pushed back under the [EU's Digital Omnibus package](#) and now phase in from late 2027, so the near-term change is mainly about transparency. The larger shift for liability is the revised [Product Liability Directive](#), which must pass into national law by 9 December 2026. It extends strict, no-fault product liability to software and AI, and ties a product's 'defectiveness' to compliance with rules like the AI Act. It applies extraterritorially, reaching UK and US firms whose AI-enabled products or services touch EU customers.


For underwriters, this blurs the line between cyber, tech E&O, and product liability, and some European carriers are already responding with 'AI security riders' and AI-specific controls, such as evidence of adversarial red-teaming. The practical question is how an AI claim gets triggered and allocated across cyber, E&O, and product wordings.

However you cut it, you need visibility

Step back from the disagreements and the same conclusion keeps surfacing. Whether AI ends up as its own line or stays spread across the ones you already write, and however you choose to cover or exclude it, you cannot manage what you cannot see. The market told us as much: the biggest gap is visibility, and better visibility would let underwriters price more accurately and write more of this risk with confidence.

The encouraging part is that visibility is achievable. Much of this exposure is externally observable. You do not need to be inside a business to tell which AI services it is running.

Most popular genAI apps based on the % of organisations using those applications



Source: [*Netskope, Cloud and Threat Report: 2026*](#)

There is also a behavioural dimension to this. With AI evolving so rapidly, how quickly an organisation identifies and addresses a problem can reveal more about its cyber resilience than the particular vulnerability, or CVE (Common Vulnerabilities and Exposures), involved. As one executive at a US-based MGA put it, **‘The fact that they knew they had it and did something about it is more important than the specific CVE.’**

Then there is the harder question of systemic risk, which silent AI already points to. The scenario one specialist said worries him most is further off but already visible: AI systems that spin up other AI systems, unknown even to the organisations running them. None of this is a reason to panic now. But you cannot judge a risk you have chosen not to look at.

The view

We have kept our own opinion out of the report until now, because the market deserves to be heard on its own terms. But we were asked the same question we put to everyone else, so here is where we land.

The separate-line debate matters less than it looks. Whether AI ends up as its own product or stays distributed across cyber, professional indemnity and general liability, the thing that decides whether it can be written profitably is the same in either case: can you see where AI is actually being used across your book? That question has to be answered at two levels at once. Organisation by organisation, because any single insured can be exposed by the AI it relies on. And across the portfolio, because the same few models sit behind so many insureds that a single failure can surface everywhere at once. The first is an underwriting problem, the second an accumulation one, and both are managed the same way: with confirmed, current exposure data, not with assumptions.

The market will settle, the way cyber eventually did. The systemic exposure may or may not be hiding in your book today. The mistake is waiting until it announces itself to find out.

“

Everyone will spend years arguing which line AI belongs in. That matters far less than whether you can see the exposure, which is already on the books, undeclared and concentrated in a few shared models. You don't need to settle the category to measure it; you need visibility. The advantage goes to whoever reads the territory.

Aaron Aanenson

Head of Insurance, North America,


KNOW YOUR EXPOSURE

**See where AI is already running
across your book – from the outside.**

Book a demo