CHUBB®

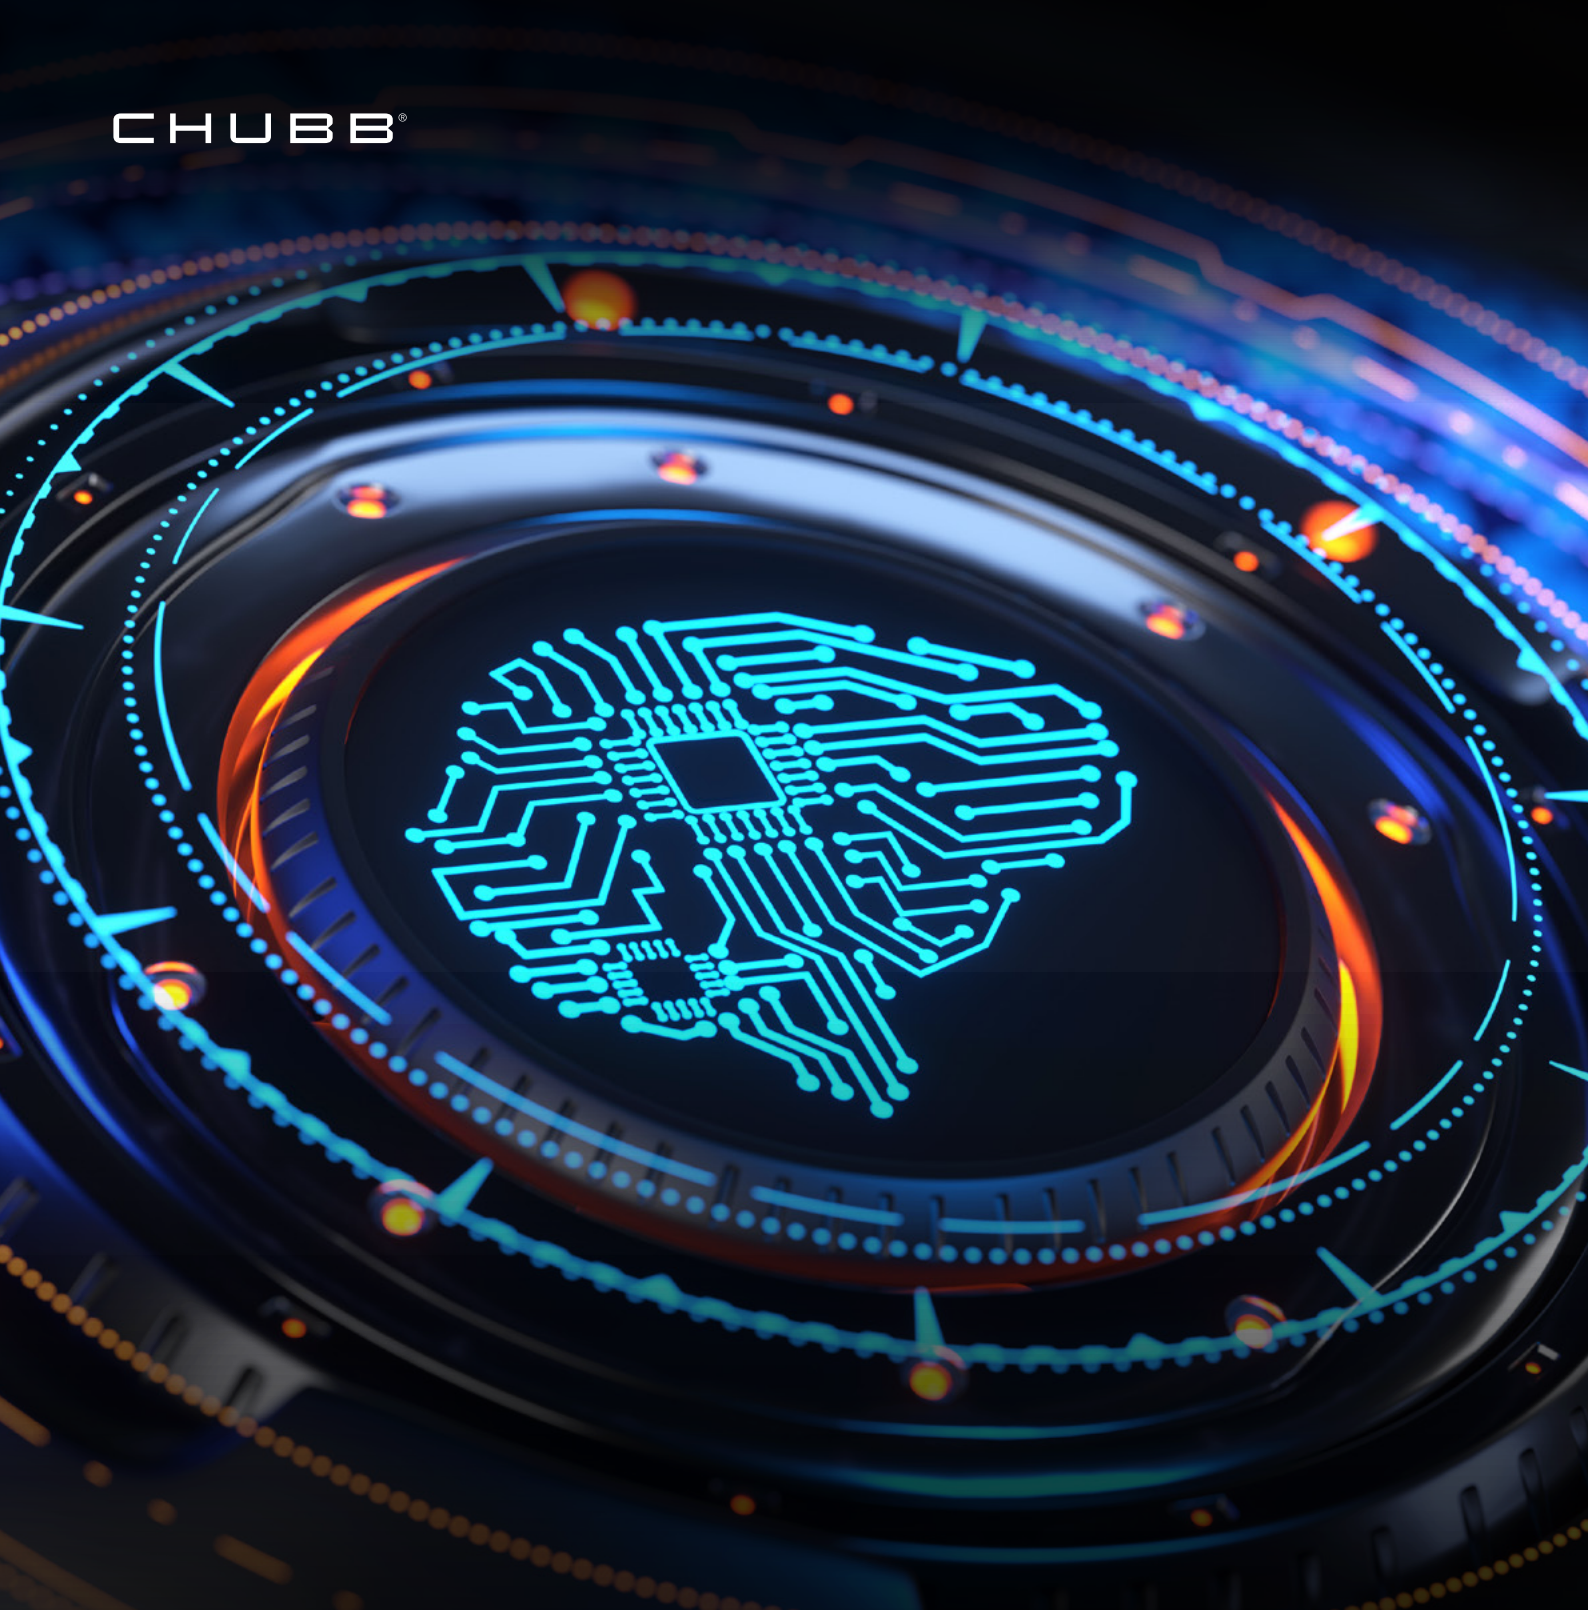

Insurance purchasing in the age of AI

How risk is changing across Cyber, Crime, D&O and Liability
... and what to do about it

Executive summary

Artificial Intelligence (AI) is increasingly part of everyday lives and almost every business process. It is frequently embedded into operational workflows across customer service, finance, HR, security, marketing and decision-making processes as well as being integrated directly into end products. For insurance purchasers, this matters because AI does not simply introduce a new standalone category of risk; it significantly reshapes existing insured perils, particularly across Cyber, Crime, Directors & Officers (D&O) and Liability/E&O.

AI changes loss outcomes through a combination of speed (incidents unfold faster than response), scale (errors replicate across transactions), opacity (harder to explain why systems acted as they did) and dependency (greater reliance on vendor supply chains). This creates a landscape where losses may become more correlated and systemic, where defence costs may rise due to evidentiary complexity and liability narratives increasingly focus on governance, oversight and organisational responsibility.

The practical outcome for Risk Managers is that many AI incidents will present as familiar claims, fraud, privacy breaches, negligent advice, defamation, discrimination or bodily injury, rather than "AI claims". This increases both severity risk and coverage complexity across towers.

To respond effectively before renewal, buyers should:

- document AI governance and ownership
- strengthen vendor chain diligence and contracting
- modernise fraud controls for deepfake-enabled threats
- run claim pathway tabletop scenarios across policies
- review policy wordings for triggered gaps or exclusion drift

Why AI is an insurance issue (and not an IT debate)

AI systems are increasingly used to make recommendations, triage cases, approve transactions, interact with customers and generate information that humans then rely upon. As AI is woven into everyday workflows, it starts to influence insured outcomes in ways that are directly relevant to buyers, brokers and underwriters.

The key point is that most AI risks do not appear as a new class of "AI claims". Instead, AI changes how traditional losses occur: it accelerates and amplifies them and makes them harder to investigate. A payroll fraud enabled by a deepfake voice is still a crime loss.

A chatbot leaking personal information is still a cyber/privacy incident. A model hallucinating incorrect advice may still be an E&O loss. But AI reshapes the facts that underpin each claim and therefore influences both liability and insurance response.

For that reason, AI should be treated as a multiplier of existing insured perils rather than a separate niche. The insurance purchasing challenge is to understand what has changed in incident mechanics and the evidentiary environment and what that means for programme design and the underwriting process.

> "AI should be treated as a multiplier of existing insured perils rather than a separate niche"

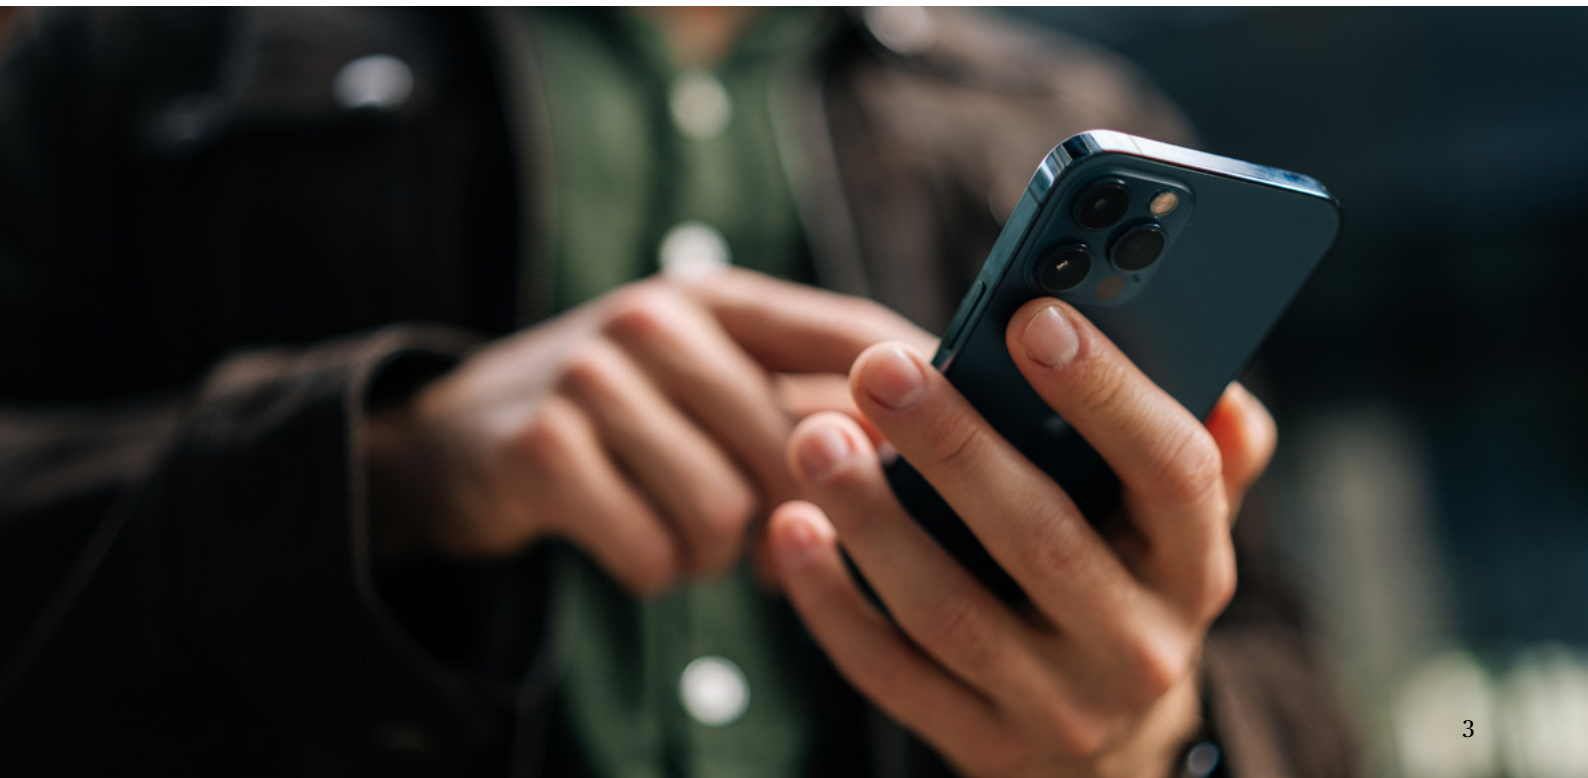

How AI changes insured loss events: six mechanisms

AI affects insured loss scenarios in a consistent set of recurring ways. These mechanisms are useful because they apply across cyber, crime, liability and governance losses and can be used as a practical framework in renewal discussions.

1

Mechanism 1 - **Speed:** incidents unfold faster than response cycles

AI can act in milliseconds and at machine scale. This means errors, exploitation and harmful outputs can propagate before an organisation detects the problem, investigates or implements a halt mechanism. This compresses response windows and increases the probability that losses escalate before containment.

2

Mechanism 2 - **Scale:** one error is replicated across many decisions

AI systems are designed to replicate behaviour at volume. If an AI system is wrong, biased or compromised, it may repeat the same mistake across hundreds or thousands of transactions, customers or operational decisions. This increases the likelihood of mass harm events rather than isolated one-off claims.

3

Mechanism 3 - **Opacity:** hard-to-explain behaviour increases defence costs

Many AI models do not provide a clear, human-readable reasoning chain. Even where outputs can be tested, it may be difficult to demonstrate exactly why a system reached a decision or whether an alternative output would have been produced under slightly different conditions. This increases evidentiary uncertainty and can raise legal costs, particularly where a claim depends on showing a defect, negligence or a failure of reasonable care.

4

Mechanism 4 - Dependency: AI introduces vendor chain risk

Unlike traditional systems built and maintained internally, AI systems often depend on multiple third parties: model providers, datasets, cloud platforms, integrators and ongoing updates. A failure or compromise may originate in any part of this chain, well outside the buyer's direct control. This creates attribution challenges and introduces exposure that spreads across vendor relationships.

5

Mechanism 5 - Manipulation: attackers can distort what the AI relies upon

AI systems can be attacked not only through classic IT compromise but also through manipulation of the information they rely on: inputs, prompts, documents, retrieval sources or training-related data. In practical terms, this can cause the model to behave as if false information is true, producing confident but harmful outputs. This expands fraud pathways and introduces new routes to consumer harm and operational breakdown.

6

Mechanism 6 - Agency: AI outputs are treated as if from a trusted representative

As AI systems make recommendations, draft communications, approve actions or guide decision-making, they increasingly function as if they were a delegated employee or agent. In claims narratives, organisations may be treated as responsible for harmful outputs because the AI system was authorised and relied upon in the workflow even where the AI vendor played a role in the failure.

Taken together, these mechanisms increase severity, correlation and uncertainty – three factors that materially change insured outcomes and can stress policy design.

The core buyer challenge: “silent AI risk” and coverage gaps

AI is best understood as a risk amplifier that will often appear inside familiar claim types. This is a critical purchasing insight: the coverage issues are rarely framed as “AI problems”. They are framed as disputes about which policy responds and whether the loss falls inside an existing definition.

Many AI incidents may trigger multiple policies simultaneously, including:

- Cyber (privacy event, network security failure)
- Crime (fraudulent instruction/impersonation)
- Tech E&O/PI (negligent advice/reliance loss)
- Media (defamation)
- D&O (oversight failures, disclosure, “AI washing”, regulatory investigations)
- General liability/consumer harm

Across the market, policy wordings vary considerably in how they address AI-related loss scenarios, which may be a consequence of differing levels of carrier experience and the speed at which underwriting language has evolved to keep pace with AI risk.

Buyers should pay close attention to:

- The scope of a covered “security failure” where AI is a contributing cause
- Whether an AI-generated error falls within a “professional services” definition
- Whether a synthetic or AI-generated output constitutes an “instruction” under crime policy language
- How exclusions apply (for example, algorithmic exclusions, professional services exclusions)
- Whether a third-party vendor’s failure is treated as “outsourced provider risk”

Accordingly, Risk Managers should assume that AI has the potential to introduce some form of coverage ambiguity and undertake proactive steps when designing and purchasing their insurance programme.

AI incident coverage mapping: how common scenarios land across policies

The table below illustrates how typical AI-related incidents are likely to be treated across insurance policies. This is indicative only, actual outcomes depend on wording, definitions, exclusions and sub-limits. Many AI losses trigger multiple lines simultaneously.

AI incident/ loss scenario	Primary loss type	Likely primary policy response	Common secondary/ overlapping policy	Key tension/ dispute points
Employee pastes PII/ PHI/trade secrets into public AI tool and data leaks	Privacy breach + regulatory	Cyber (privacy, breach response, regulatory defence)	Tech E&O (if client harm), D&O (if governance)	Was there a failure by the insured to adhere to data handling requirements? Did the breach violate applicable privacy regulations?
Prompt injection attack causes chatbot to disclose sensitive customer info	Privacy breach + security event	Cyber	Tech E&O	Does the policy cover "Computer Malicious Acts" and provide privacy liability coverage?
AI system hallucinates and gives incorrect instructions causing customer financial loss	Professional negligence/ misrepresentation	Tech E&O/ Professional Indemnity	GL (rare), Media	Whether "professional services" exclusion applies; causation and reliance proof
AI chatbot defames a person or outputs discriminatory content	Third-party content liability	Media liability (or Tech E&O if bundled)	GL (sometimes), EPL (if employee context)	Publication/defamation triggers, intent exclusions, content exclusions in E&O
AI-generated marketing makes false claims (performance, pricing, "guaranteed outcomes")	Consumer protection/ deceptive trade	Tech E&O/Media	D&O (if investor impact)	Misrepresentation vs deliberate deception; intentional wrongdoing exclusions
AI creates copyrighted content (or very similar) and IP holder sues	IP infringement defence/costs	Media liability (best fit), sometimes Tech E&O	D&O (if public company disclosure), Cyber (rare)	IP exclusions differ widely (copyright vs trademark vs patent); sub-limits often apply
AI trained on allegedly stolen trade secrets/data	IP + trade secrets	Tech E&O (if written broadly)	D&O	Policies often exclude IP/trade secret claims unless negotiated
AI model used in HR creates discriminatory hiring outcomes	Employment discrimination	EPL	D&O (oversight), Tech E&O (vendor risk)	Whether automated decisioning violates laws; coverage for regulatory actions varies
AI credit/ underwriting/pricing creates unfair/ discriminatory outcomes	Regulatory + class action	Tech E&O/ GL (varies), sometimes D&O	D&O	Some policies exclude pricing decisions; regulators may pursue non-insurable penalties

AI incident/ loss scenario	Primary loss type	Likely primary policy response	Common secondary/ overlapping policy	Key tension/ dispute points
AI-driven denial/ triage in claims or benefits causes systemic harm	Operational + consumer class action	Tech E&O/ Professional Indemnity	D&O	Negligence, unfair practices, failure of oversight; defence costs may escalate significantly
Deepfake impersonation causes fraudulent funds transfer	Financial loss	Crime (social engineering/funds transfer fraud)	Cyber (if system intrusion), D&O (oversight)	Whether "instruction" includes voice/video; voluntary parting and authentication controls
AI-driven invoice/ payment fraud with spoofed vendor details	Financial loss	Crime	Cyber	Disputes whether "computer fraud" applies; social engineering sub limits
AI tool used in finance auto-generates payments with wrong beneficiary	Operational error	Crime (depends), or E&O	D&O	Was it "fraud" or "error"? Many crime policies do not cover pure error
Ransomware extortion accelerated by AI-based pre-attack reconnaissance	Cyber extortion	Cyber	Crime (rare)	Systemic risk exclusions and vendor-related disputes
AI vulnerability in vendor tool leads to security event	Cyber + vendor	Cyber	Tech E&O	Third-party failure and contingent business interruption triggers
AI model drift causes incorrect outputs leading to customer losses	E&O loss	Tech E&O	D&O	Whether drift monitoring was "reasonable"; may become a governance dispute
AI incident triggers regulator investigation into governance and controls	Investigation costs	D&O (investigation coverage varies)	Cyber (if privacy/ security)	Whether investigation constitutes a "claim"; insurability of penalties
Public company stock drops after AI failure controversy	Securities class action	D&O	—	Disclosure, materiality, reliance, scienter defences
External AI assurance provider certifies model which later causes harm	Reliance-based negligence	Tester's PI/Tech E&O (vendor's policy)	Buyer's Tech E&O/D&O	Contractual limitation battles, third-party reliance arguments
AI outages disrupt operations (vendor outage)	Business interruption	Cyber (if tech outage coverage exists)	Property BI (rare), Supply chain	Whether "system failure" coverage exists, and if vendor failure is included

Buyer action tip: require your broker to provide a “single-incident multi-policy response view” for at least three AI scenarios. Many disputes arise not from lack of insurance, but from uncertainty about which policy is primary.

Practical risk management by insurance line

The following guidance translates the six AI mechanisms into line-of-business specific risks and practical steps. Risk Managers should use these as a working checklist alongside their renewal process.

Cyber: AI expands the attack surface beyond IT

AI introduces new data leakage routes (employees entering sensitive information into AI tools), new exploitation techniques (prompt injection) and new vectors for harmful customer-facing outputs. AI adoption also increases dependency on third-party AI platforms and cloud services, which may complicate incident ownership and reporting.

What Risk Managers should do:

- Implement allow-lists for AI tools and DLP controls
- Log and monitor AI tool usage and outputs where feasible
- Conduct vendor diligence on model providers, hosting and subcontractors
- Confirm policy response for privacy events involving AI services

D&O: oversight, disclosure and governance become central

As AI becomes operationally material, failures and controversies become governance-level events. Exposure may arise from allegations of inadequate oversight, misleading disclosures, consumer harm, discriminatory outcomes or misrepresentation about AI capability and compliance.

What Risk Managers should do:

- Treat governance documentation as a defence asset
- Ensure clarity on investigations, regulatory enquiries and claim triggers
- Stress-test Side A adequacy and entity coverage
- Audit public statements to reduce "AI washing" allegations

Crime: impersonation fraud becomes scalable and credible

Deepfake voice/video and high-quality language imitation reduce the reliability of human intuition as a fraud control. AI-supported impersonation can enable fraudulent payment approvals, vendor invoice redirection and social engineering at scale.

What Risk Managers should do:

- Remove authentication reliance on voice/video alone
- Strengthen out-of-band verification for sensitive approvals
- Test whether "fraudulent instruction" wording captures synthetic impersonation scenarios
- Ensure operational controls align with expected insurer position

Liability/Tech E&O/Media: reliance harms, incorrect outputs, IP and reputational loss

AI-generated outputs can cause downstream losses when relied upon by customers, staff or counterparties. These include hallucinated advice, incorrect triage decisions, reputational harm, defamation and IP disputes.

What Risk Managers should do:

- Place guardrails on customer-facing AI outputs
- Clarify E&O/media overlaps and exclusions
- Contractually manage vendor indemnities and responsibility allocation
- Carefully evaluate and manage AI project responsibilities, liabilities and indemnifications
- Ensure testing and developer responsibilities are clearly defined between technology provider and client

Key Questions for brokers and underwriters on AI

The following questions are designed to elicit concrete underwriting positions and reveal silent gaps. Risk Managers should raise these as a structured agenda item at renewal.

Programme structure and silent AI exposure

- Where do you expect our largest AI-driven losses to land: Cyber, Crime, D&O, EPL, Tech E&O, Media or GL?
- Do any policies in our tower contain AI/algorithm/model-related exclusions or limitations? Please list them verbatim and explain impact.
- Are any sub-limits likely to apply to AI-shaped claims (social engineering, privacy regulatory, media/IP, tech errors)?

Cyber-specific

- Does our cyber policy treat AI prompt leakage or chatbot prompt injection as a covered privacy event? What wording supports this?
- Does the policy cover data privacy incidents arising from employee use of third-party AI tools (copilots, LLM platforms)?
- Is vendor AI outage included within business interruption/dependent business interruption coverage?

Crime and social engineering

- Does "fraudulent instruction" explicitly include voice/video deepfake impersonation, or only written/email instruction?
- Are there exclusions (voluntary parting/authorised instruction/authentication failure) likely to defeat deepfake claims?
- What controls does the underwriter expect us to implement (call-back, dual approval, out-of-band verification) and will these be warranties?

D&O/ governance

- Does the D&O policy cover pre-claim regulatory inquiries linked to AI governance (data use, discrimination, AI safety)? What triggers?
- How does the insurer view AI as a "mission critical risk" for oversight purposes, and what board-level governance evidence will improve terms?
- Does the insurer have a view on "AI washing" disclosure claims, and do they expect specific AI-related disclosures in risk factors?

Liability/E&O/Media

- What policy responds to AI hallucinations and incorrect advice outputs that cause customer loss; Tech E&O, PI, or GL? Where are the gaps?
- How are IP claims treated when the harm is AI-generated (copyright/trademark/patent)? Are defence costs covered? Are there sub-limits?
- Do policies exclude claims "arising out of" algorithms or automated decisioning, even if AI is only one contributing cause?

Conclusion

AI is reshaping traditional insured losses by increasing speed, scale, dependency and evidentiary uncertainty. For insurance buyers, the primary challenge is not simply greater incident likelihood it is greater severity, correlation and uncertainty in coverage coordination across towers. The most effective response is governance-driven and practical: document material AI uses, strengthen vendor chain diligence, modernise fraud controls, conduct scenario-based coverage mapping and clarify wordings before exclusions and claims precedents harden.

The regulatory landscape adds a further dimension that will require ongoing attention and is likely to become increasingly consequential for insurance buyers. In the EU, structured governance duties under the AI Act are already taking effect for higher-risk uses, while product liability concepts are expanding to encompass AI-enabled functionality and software.

In the US, a patchwork of NIST guidance, sector regulator activity, state-level rules and litigation-driven accountability is generating significant and rapidly evolving exposure. The practical message across both jurisdictions is consistent: defensibility depends on demonstrating reasonable governance and oversight, not technical sophistication alone.

Beyond the EU and US, jurisdictions across Asia-Pacific, the Middle East and Latin America are at different and accelerating stages of developing AI-specific frameworks, sector-level guidance and liability standards.

The shape of AI-related regulatory exposure across these markets and what it means for multinational programme design, will be a subject for further examination as this landscape matures.

It is an area that warrants close and sustained attention from Risk Managers and their brokers as renewal conversations evolve.

CHUBB®

9184-HS 04/26