



WHITE PAPER

Your Agents Are Ready. Is Your Insurer?

Why Agentic AI's Deployment Gap Is more likely to be
an Insurance Problem than a Technology One

June 2026

By Dr. Raphael Troitzsch, Nadine Moore and Coralie Ming

Your Agents Are Ready. Is Your Insurer?

Why Agentic AI's Deployment Gap Is more likely to be an Insurance Problem than a Technology One

In 2024, a fleet of oil tankers lay anchored in the Gulf of Oman. The cargo was ready. Military escort had been arranged. The route through the Strait of Hormuz was physically open, but the tankers did not move. Not because of the geopolitical risk, which was understood, nor the military calculus. But because no insurer would write the policy in its original price range¹. The economic consequences of that judgment—cargo immobilized, supply chains disrupted, revenues vaporized—exceeded the cost of several medium-sized military engagements. The underwriting decision produced an economic consequence that military presence alone could not prevent.

Across the Pacific, autonomous vehicle taxis in California have cleared every engineering and regulatory hurdle set by developers. Safety records, mapping fidelity, millions of operational miles. Yet dozens of hotels in the state still bar them from their driveways. Not by statute or by engineering concern, but because of an insurance clause: the property policy excludes autonomous vehicle liability, no carve-out has been negotiated, and the underwriter has not been asked to consider one². The technology works. The coverage doesn't.

These are not edge cases. They are previews.

The constraint that will most visibly determine the pace of the agentic AI economy over the next decade is not the model, not the chip, not the data center, and not—despite the volume of ink spent on the subject—the regulation. It is the insurance policy. More precisely, the absence of one.



¹ Safety concerns, not insurance availability, driving reduced vessel traffic in the Strait of Hormuz. Lloyd's Market Association, 2026. <https://lmalloyds.com/safety-concerns-not-insurance-availability-driving-reduced-vessel-traffic-in-the-strait-of-hormuz/>

² Autonomous Vehicle Insurance Policy Exclusion Survey. Insurance Information Institute, 2025. <https://www.iii.org/article/background-on-self-driving-cars-and-insurance>

The Gap is Structural, not Transitional

The insurance industry's response to agentic AI has been widely characterized as a lag—the market will catch up, premiums will find their level, products will emerge. This characterization is wrong, and the error is consequential. What the market is experiencing is not a lag, but a structural exclusion, and that distinction matters enormously for anyone planning three-to-five-year deployments now.

As of January 2026, a standard Insurance Service Office (ISO) endorsement—the template language for the majority of US commercial property and casualty policies—contains a blanket technology-neutral exclusion for AI systems across business interruption and products liability coverage³. It carries no transition language. No carve-outs for mature or certified systems. No timeline for review. It simply excludes the risk class. That endorsement now appears in roughly 82% of commercial P&C templates across the US market⁴. Insurers did not draft it carelessly. They drafted it because they had to.

The challenge here is not institutional inertia or lack of commercial intent. It is a genuine data architecture problem. Commercial insurance pricing depends on historical loss series. Regulators require auditable actuarial models. Actuarial models require loss history. And loss history, for agentic AI systems operating across enterprise environments, does not yet exist in any statistically meaningful form⁵. The industry faces a risk it has not previously encountered: one that has never failed at scale, in conditions that cannot yet be fully specified, with cascades that cannot reliably be bounded in advance. There is simply no past to draw from.

The result is a coverage gap that dwarfs anything the industry has faced in modern memory. Genuine deployment-blocked exposure—the segment of enterprise AI spend that cannot proceed without coverage for operational failure—sits at an estimated \$150 to \$200 billion annually^{6,7}. Against that figure, the current aggregate capacity of AI-specific specialty insurance products is approximately \$25 million⁸. Demand exceeds capacity by a ratio of between 6,000 and 8,000 to one. For context, when the automotive, aviation, nuclear energy, and cyber industries each confronted their formative insurance crises, the equivalent gap never exceeded a ratio of 100 to one^{9,10,11,12}. The agentic AI gap is between 60 and 80 times larger than any prior technology has produced at a comparable deployment stage.

This is not a coverage delay. It is a coverage architecture engineered to be stable.

³ ISO Commercial General Liability Endorsement CG 40 47. ISO Properties, Inc., 2023. <https://www.verisk.com/products/forms-rules-and-loss-costs/>

⁴ AI Insurance Market Coverage Survey. Marsh McLennan, 2025. <https://www.marshmma.com/us/insights/details/business-insurance-trends.html>

⁵ Actuarial Seasoning Requirements—Loss Data Standards. Casualty Actuarial Society, 2023. <https://www.casact.org/sites/default/files/2021-05/Statement-Of-Principles-Ratemaking.pdf>

⁶ AI Insurance Market Coverage Survey. Marsh McLennan, 2025. <https://www.marshmma.com/us/insights/details/business-insurance-trends.html>

⁷ Specialty AI Insurance Market Capacity Report. Lloyd's Market Association, 2025. <https://lmalloyds.com/lma-survey-maps-underwriters-views-of-ai-loss-scenarios-across-key-lines/>

⁸ Specialty AI Insurance Market Capacity Report. Lloyd's Market Association, 2025. <https://lmalloyds.com/lma-survey-maps-underwriters-views-of-ai-loss-scenarios-across-key-lines/>

⁹ Price-Anderson Nuclear Industries Indemnity Act—Legislative History. U.S. NRC, 1957/2005. <https://www.nrc.gov/reading-rm/doc-collections/nuregs/staff/sr0980/>

¹⁰ TRIA: Terrorism Risk Insurance Program Reauthorization Act. U.S. Treasury, 2002/2022. <https://home.treasury.gov/policy-issues/financial-markets-financial-institutions-and-fiscal-service/federal-insurance-office/terrorism-risk-insurance-program>

¹¹ Pool Re Annual Report 2024. Pool Reinsurance Company, 2024. <https://www.poolre.co.uk/resources/annual-reports/>

¹² London Market Cyber Capacity Study. Lloyd's of London, 2024. <https://www.lloyds.com/about-lloyds/our-market/what-we-insure/cyber/cyber-risk-insight>

Three Lines, Three Forms of Brake

The exclusion problem manifests differently across the three insurance lines most material to agentic deployment: business interruption, cyber, and liability. Each creates its own form of drag.

Business interruption coverage, as currently structured across most commercial policies, requires physical damage to trigger a valid claim: a fire, a flood, a tangible event¹³. An AI agent that halts a supply chain by misrouting logistics decisions, that freezes a financial institution's trading operations through a model error, or that disrupts a healthcare system's patient management workflow produces no physical damage. It produces an outage. Under existing BI policy language, that outage is largely uninsured, regardless of its economic impact. As enterprise environments become increasingly agentic and more workflows depend on AI systems making decisions in real time, the category of economically devastating outages that fall outside BI coverage also grows larger. The coverage architecture was built for industrial infrastructure, not for software that acts.

Cyber insurance presents a second, more technically complex problem: concentration risk at a scale the industry has not previously been asked to hold. Enterprise agentic deployments overwhelmingly depend on a small number of foundation model providers, cloud infrastructure platforms, and identity management systems¹⁴. This concentration produces correlated failure modes. When a core infrastructure component fails—whether through a zero-day vulnerability, a model weight error, or a cloud platform outage—it does not affect just one policyholder. It fails for thousands simultaneously. The statistical correlation of AI failure events across enterprise deployments is estimated at approximately 0.72—high enough to fundamentally break the pooling mathematics on



¹³ Business Interruption Insurance—Physical Damage Trigger Analysis. Swiss Re Institute, 2024. <https://www.swissre.com/dam/jcr:f6a2a603-8f7c-483e-bab3-8dd9c7aa52e4/swiss-re-institute-expertise-publication-complex-supply-chains.pdf>

¹⁴ Foundation Model Concentration Risk Report. Financial Stability Board, 2024. <https://www.fsb.org/2024/11/fsb-assesses-the-financial-stability-implications-of-artificial-intelligence/>

which standard insurance pricing depends¹⁵. A loss model built on the assumption of independent events produces inaccurate risk estimates when the underlying events are highly correlated. For underwriters building books of AI-related coverage, correlation adjustment is not a refinement; it is a foundational requirement of sound AI risk pricing. The cyber insurance market learned this the hard way: early movers who built broad, uncorrelated books before the concentration dynamic was properly understood absorbed losses that fundamentally reshaped the market's appetite and led to significant portfolio corrections across several major carriers¹⁶.

Liability presents the third dimension: the question of legal responsibility when an autonomous system causes harm¹⁷. Auto insurance, construction liability, professional indemnity, product liability—each of these lines was designed around the implicit assumption of a human decision-maker in the chain. A driver. A contractor. A doctor. A manufacturer. Agentic AI systems disrupt that structural idea of liability. When an AI agent makes the wrong medical recommendation, executes an erroneous financial trade, or causes a physical accident while operating autonomously, the liability attribution question—whether responsibility lies with the developer, deploying enterprise, or individual who initiated the task—is legally unresolved and remains unaddressed in most policy language. In the absence of coverage, enterprises do not accept this liability. They are not deploying the system.



¹⁵ Cloud Provider Dependency and Correlated Failure Risk in Enterprise AI. NIST, 2024. <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>

¹⁶ London Market Cyber Capacity Study. Lloyd's of London, 2024. <https://www.lloyds.com/about-lloyds/our-market/what-we-insure/cyber/cyber-risk-insight>

¹⁷ AI Liability Attribution: Legal Framework Analysis. OECD, 2024. <https://legalinstruments.oecd.org/api/print?id=648&lang=en>

Agents are not Utilities

The analogy most frequently invoked by those who believe the coverage gap will resolve itself is the electricity grid. We insure the grid, which is also complex, the argument runs. Therefore, we will learn to insure AI. The failure of this analogy illuminates the depth of the structural problem.

Power grids break down in well-characterized ways. The failure modes—overload, equipment fault, weather event, cyberattack on SCADA systems—are enumerated, modeled, and have generated decades of loss history. Underwriters know what a grid outage looks like, how far it propagates, and what the economic impact has historically been. The data exists. The models are defensible.

Agentic AI systems have none of these properties, and the contrast runs deeper than the absence of historical data. Failure modes are not enumerated—an agent can fail by producing subtly incorrect outputs over a sustained period, by correctly executing incorrectly specified instructions, or through emergent interactions between multiple agents that no individual system was designed to produce. The impact radius is poorly bounded, the recovery time unclear, and the failure taxonomy of a model deployed today will not be the same failure taxonomy of its successor in eighteen months.

The Velocity Mismatch

Actuarial seasoning—the accumulation of loss data sufficient to build a defensible pricing model—requires five to seven years of claims experience under stable conditions¹⁸. The half-life of a deployed AI model is eighteen to twenty-four months¹⁹. The actuarial clock runs 2.5 to nearly 5 times slower than the technology clock. The insurance industry cannot organically accumulate the data it needs to price the risk because the risk keeps changing faster than the data can be observed and analyzed. This is not a problem that resolves with patience; it is a structural velocity mismatch, and no amount of market maturation changes the underlying arithmetic.

¹⁸ Actuarial Seasoning Requirements—Loss Data Standards. Casualty Actuarial Society, 2023. <https://www.casact.org/sites/default/files/2021-05/Statement-Of-Principles-Ratemaking.pdf>

¹⁹ Agentic AI Deployment Survey: Enterprise Risk Readiness. Gartner, Inc., 2025. <https://www.gartner.com/en/newsroom/press-releases/2025-09-30-gartner-survey-finds-just-15-percent-of-it-application-leaders-are-considering-piloting-or-deploying-fully-autonomous-ai-agents>

The Economic Drag is already Measurable

Gartner forecasts that more than 40% of agentic AI projects will be cancelled by 2027. This is striking not for its lack of precision—forecasts of this type never are—but for its direction²⁰. The gap between planning and execution is not uniform across organizational types, and understanding its shape matters for how the industry responds.

For large enterprises, the coverage gap operates as a hard gate. Legal and risk teams encountering exclusions, unresolved liability language, and absent policy language halt procurement, delay vendor agreements, and in many cases kill projects that have otherwise cleared every hurdle²¹. The technology exists, the business case has been made, and the insurance desk says no. Large foundation model providers and platform operators sit somewhat apart from this dynamic: having constructed private risk transfer mechanisms and carrying sufficient balance sheet to absorb meaningful operational exposure, they are insulated in ways that most enterprises are not.

For mid-market organizations, smaller enterprises, and startups, the dynamic is subtler and in some ways more consequential. These organizations frequently deploy agentic AI without conducting any insurance gap assessment—not because they have resolved the underlying risk, but because they lack the institutional infrastructure to identify it in the first place. They accept the exposure implicitly, without the ability to quantify what a serious incident would actually cost. This is a different problem from the procurement blockage that slows large enterprises. It is unquantified tail risk, carried by organizations that in



²⁰ Agentic AI Deployment Survey: Enterprise Risk Readiness. Gartner, Inc., 2025. <https://www.gartner.com/en/newsroom/press-releases/2025-09-30-gartner-survey-finds-just-15-percent-of-it-application-leaders-are-considering-piloting-or-deploying-fully-autonomous-ai-agents>

²¹ AI Insurance Market Coverage Survey. Marsh McLennan, 2025. <https://www.marshmma.com/us/insights/details/business-insurance-trends.html>



many cases could not survive a material AI-caused incident.

That risk materialized in April 2026, when a Cursor AI agent assigned to fix a routine staging bug at PocketOS—a SaaS startup serving car rental operators—located an over-privileged API token within the same session and used it to delete the company's entire production database and all volume-level backups²². The process took nine seconds. The agent subsequently logged that it had "violated every principle it was given." PocketOS survived only because Railway's CEO happened to learn of the incident and personally restored the data within thirty minutes. Without that intervention, the oldest recoverable backup would have been three months old, and there was no insurance coverage in place to fill the gap. The incident captured, at operational speed, the exposure that uninsured agentic deployment carries for smaller organizations: a single autonomous action can produce consequences that no recovery plan can outrun.

The aggregate productivity cost of the coverage gap—across the deployment blockage at large enterprises and the uninsured exposure at smaller ones—runs to tens of billions of dollars annually²³. It compounds as timelines extend and agentic systems become more deeply embedded in critical workflows, and it concentrates sharply with industries that are more dependent on agentic systems to address structural labor and efficiency challenges—the very industries where this technology potentially offers the most value.

²² "Cursor-Opus agent snuffs out startup's production database." The Register, April 27, 2026. https://www.theregister.com/2026/04/27/cursoropus_agent_snuffs_out_pocketos/

²³ AI Economic Impact and Productivity Multipliers. BCG Henderson Institute, 2025. <https://www.bcg.com/publications/2025/are-you-generating-value-from-ai-the-widening-gap>

What A Credible Path Forward Requires

The historical record offers a guide, if not a template. Every industry that has navigated a structural insurance gap of this magnitude—nuclear energy in the late 1950s, post-9/11 terrorism risk financing, flood risk in the United Kingdom through Flood Re, and the emergence of the London market's specialist cyber capacity in the 2000s—has done so through deliberate architecture rather than organic market development^{24,25,26,27,28}. In each case, the gap closed not because underwriters grew more comfortable with the risk over time, but because the market structure was redesigned to make the risk insurable. Actuarial data was engineered into existence through mandated reporting and disclosure frameworks. Capacity was created through consortium structures that allowed individual carriers to hold bounded exposure. Sovereign or quasi-sovereign backstops were deployed at the tail—not at the primary layer, where they would suppress private capacity formation, but at the catastrophic layer, where private capital cannot rationally go without government participation.

The Flood Re Design Principle

Launched in 2016, Flood Re created a reinsurance pool funded by a market-wide levy—distributing the cost of flood risk across the entire home insurance market rather than concentrating it on exposed policyholders or individual carriers²⁹. Critically, it was architected from the outset to be temporary: a sunset provision mandates full transition back to private market pricing by 2039. That sunset is not cosmetic. It is the mechanism that preserves private carriers' incentive to continue pricing and underwriting, preventing the pool from becoming a permanent subsidy that crowds out commercial capacity. An AI risk architecture that borrows this design—a levy-funded transitional pool with bounded sovereign exposure at the catastrophic tail and a clear sunset horizon—could replicate Flood Re's discipline at the scale the agentic economy requires.

²⁴ Price-Anderson Nuclear Industries Indemnity Act—Legislative History. U.S. NRC, 1957/2005. <https://www.nrc.gov/reading-rm/doc-collections/nuregs/staff/sr0980/>

²⁵ TRIA: Terrorism Risk Insurance Program Reauthorization Act. U.S. Treasury, 2002/2022. <https://home.treasury.gov/policy-issues/financial-markets-financial-institutions-and-fiscal-service/federal-insurance-office/terrorism-risk-insurance-program>

²⁶ Pool Re Annual Report 2024. Pool Reinsurance Company, 2024. <https://www.poolre.co.uk/resources/annual-reports/>

²⁷ London Market Cyber Capacity Study. Lloyd's of London, 2024. <https://www.lloyds.com/about-lloyds/our-market/what-we-insure/cyber/cyber-risk-insight>

²⁸ Flood Re Annual Report 2024. Flood Reinsurance Ltd., 2024. <https://www.floodre.co.uk/wp-content/uploads/Flood-Re-Limited-ARA-2024-EJ-Signed.pdf>

²⁹ Flood Re Annual Report 2024. Flood Reinsurance Ltd., 2024. <https://www.floodre.co.uk/wp-content/uploads/Flood-Re-Limited-ARA-2024-EJ-Signed.pdf>

The Cyber Insurance Precedent: Narrow First, Expand as Data Accumulates

When cyber risk emerged in the late 1990s, the market faced the same foundational barriers: no loss history, no shared taxonomy of what a covered event meant, and no actuarial basis for pricing systemic exposure. Some carriers declined. Others built concentrated books on inadequate data and absorbed the losses³⁰.

The market found its way not by solving the data problem, but by reframing it—starting with what it could confidently price. Breach response costs first: forensics, notification, credit monitoring—bounded, observable, and sufficiently isolated from systemic correlation. Ransomware payments with defined sublimits. Business email compromise with specific triggers. Systemic cyber coverage—the kind that pays for cascading failure across thousands of policyholders—remains largely unwritten two decades on. But the cyber market is functional and substantial precisely because it started with what it could manage.

The AI market will follow the same path. The first generation of products will be tightly scoped by design: bounded AI incident response costs, regulatory fine cover for specific compliance failures, professional liability for narrowly defined AI-assisted decisions. Explicit sublimits, specific triggers, systemic exclusions. That is not a limitation of ambition; it is the correct architecture for a market at this stage of data maturity, and the foundation on which broader coverage will be built as loss experience accumulates.

For insurers, the question is not whether to wait for AI risk to become comprehensively modelable, but which bounded risks can be priced with confidence today and how to begin writing those products while the broader data architecture matures. Carriers that invest in loss-data collection, taxonomy standards³¹, and the policy conversations that will govern AI liability will be positioned to expand as the market develops. Those that wait will find themselves writing into a market already priced by others.

For enterprise deployers, the question is whether to treat the coverage gap as a constraint or a design problem. Organizations restructuring vendor agreements, building auditable failure records into their AI deployment architecture, and engaging insurers in pre-placement dialogue³² are creating the loss-data infrastructure that will eventually support coverage. Those that deploy and hope are accumulating uninsured exposure.

The architecture that resolves this—the combination of procurement-side standards, parametric underwriting capacity, and appropriately structured sovereign backstops—is knowable. The analogs exist. The structural barriers are technical and political, not conceptual. The design work is genuine, the stakeholder alignment is complex, and the window for getting the architecture right before a significant AI-caused loss event forces a less deliberate resolution is measured in months, not years. What happened in the Gulf of Oman was a preview of what will happen when a large, consequential economic activity collides with an insurance market that was not designed for it. The agentic economy is a significantly larger activity. The collision, if unaddressed, will be correspondingly larger.

The tankers are already at anchor. The question is whether the next fleet will sail.

³⁰ London Market Cyber Capacity Study. Lloyd's of London, 2024. <https://www.lloyds.com/about-lloyds/our-market/what-we-insure/cyber/cyber-risk-insight>

³¹ Parametric Insurance: Innovations in Trigger Design. Geneva Association, 2024. <https://www.genevaassociation.org/publications/research-reports>

³² FAR Part 12 / Commercial Item Policy—AI Procurement. U.S. GSA / FAR Council, 2024. <https://www.acquisition.gov/far/>

About the Authors



Raphael Troitzsch is a Managing Director and Partner in BCG's Zürich office. You may contact him by email at troitzsch.raphael@bcg.com



Nadine Moore is a Managing Director and Senior Partner in BCG's Chicago office. You may contact her by email at moore.nadine@bcg.com



Coralie Ming is a Principal in BCG's Zürich office. You may contact her by email at ming.coralie@bcg.com

For Further Contact

If you would like to discuss this report, please contact the authors.



Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.

For information or permission to reprint, please contact BCG at permissions@bcg.com. To find the latest BCG content and register to receive e-alerts on this topic or others, please visit bcg.com. Follow Boston Consulting Group on [LinkedIn](#), [Facebook](#), and [X \(formerly Twitter\)](#).



BCG