



Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

James Tuplin, SVP and Head of International Cyber, Mosaic;
Jay Vinda, Global CISO and Cyber Risk Engineering Lead, Mosaic;
Sanith Wijesinghe, Chief Engineer, Public Sector, MITRE;
Chris Folk, Director, Cyber Solutions Cell, MITRE

“

The emergence of frontier AI models has triggered a debate across the cyber security and risk management communities – what happens when attackers gain access to systems capable of discovering vulnerabilities, generating exploits and reasoning through attack paths at unprecedented speeds. While this is a reasonable debate to be having, this framing has the risk of overstating what is actually changing.

This Thought Leadership Paper Series is a joint project from Mosaic Insurance and MITRE.

Introduction | The New Economics of Cyber Risk

How CISOs, Insurers, and Risk Managers Respond to the Age of Frontier Artificial Intelligence

Paper 1 | Beyond Patch Velocity

Rethinking Cybersecurity Strategy in a Post-Frontier AI World

Paper 2 | Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

Paper 3 | Buying Resilience

How Risk Managers Can Close the Gap Between the CISO and the Boardroom

Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

The emergence of frontier AI models has triggered conversations across the cybersecurity industry, many of which focus on a familiar issue: what happens when attackers gain access to systems capable of discovering vulnerabilities, generating exploits, and identifying attack paths at unprecedented speeds, particularly speeds that outpace human remediation efforts?

Cyber insurance has always existed in a world of uncertainty. Insurers and cybersecurity practitioners do not know when a new vulnerability will be discovered, whether it will be disclosed responsibly, how quickly it can be weaponized, or how broad the scope of affected entities will be. While this uncertainty is not new, its scale is changing. Frontier AI has dramatically increased the rate at which vulnerabilities are discovered, reduced the effort required to identify exploit paths, and compressed the timeline between discovery and exploitation. At the same time, remediation remains constrained by operational realities that cannot easily be accelerated at the same pace.

The evidence of this shift is already visible in production patch cycles across major software vendors. VulnCheck reported in May 2026 that Common Vulnerabilities and Exposures (CVE) disclosure volumes were up sharply year-to-date across several major software suppliers: Chrome up 563.2%, VMware up 180.9%, Apache up 170.3%, Mozilla up 156.9%, HPE up 132.3%, and F5 up 113.8%. GitHub CVE issuance was up 476.07% over the same period. VulnCheck notes the increases are consistent with broader use of AI-assisted vulnerability discovery, though the signal is still emerging and not all of it can be directly attributed to AI. Chrome is the clearest illustration: 194 CVE disclosures across the whole of 2025, against 378 by mid-May 2026. Mozilla's experience is more precisely documented: using Claude Mythos Preview as part of an agentic hardening pipeline, Mozilla fixed 271 Firefox security bugs, though Mozilla's own account notes that the pipeline still required significant engineering work around deduplication, triage, verification, patching, review, and release management. Cloudflare took a similar approach internally: Cloudforce One tasked an AI coding agent with a vulnerability analysis of its own systems, uncovering CVE-2026-22813, scored 9.4 on the Common Vulnerability Scoring System (CVSS) scale, a critical flaw in Markdown rendering pipelines allowing unauthenticated remote code execution. The pattern across vendors is consistent: AI-scale discovery is not producing marginal increments to existing vulnerability counts. It is producing step changes.

AI has dramatically increased the rate at which vulnerabilities are discovered, reduced the effort required to identify exploit paths, and compressed the timeline between discovery and exploitation.

Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

Frontier AI models are also demonstrating something more troubling than raw discovery volume. They are showing the ability to develop credible, working exploit paths even in the presence of standard defensive controls by chaining together vulnerabilities that no conventional prioritization model would have flagged as high risk.

The ExploitGym benchmark, published in May 2026 by researchers from UC Berkeley, Anthropic, OpenAI, and Google, tested frontier AI agents against 898 real-world vulnerabilities, toggling standard mitigations such as Address Space Layout Randomization (ASLR), stack canaries, and the V8 heap sandbox on and off to isolate their effect. With mitigations disabled, Claude Mythos Preview successfully exploited 157 of those vulnerabilities and GPT-5.5 exploited 120. With mitigations enabled, success rates dropped sharply but did not reach zero: Mythos Preview still succeeded on 45 and GPT-5.5 on 21. Critically, these results were achieved not just against known high-severity targets. Agents routinely identified exploit paths by chaining multiple lower-severity findings into sequences that achieved full system compromise, arriving at outcomes that no individual CVSS score in the chain would have predicted. Microsoft's MDASH system confirmed this pattern in production: CVE-2026-33824 was discovered by reasoning across six source files and identifying an aliasing ownership bug whose exploitability only became apparent when the flawed implementation was compared with the correct version of the same pattern elsewhere in the codebase. No single-file analysis and no conventional scanner would have found it.

This matters for the cyber insurance market in a specific way. Most organizations carry what practitioners call a graveyard: a backlog of known medium- and low-severity vulnerabilities that have been assessed, deprioritized as not immediately exploitable, and left unpatched indefinitely. That deprioritization was reasonable under a threat model in which chaining those vulnerabilities into a working exploit required specialist expertise, significant time, and a degree of architectural understanding that limited the pool of capable adversaries. Frontier AI collapses that assumption. Chaining was already trending longer before AI entered the picture. Google Threat Intelligence Group's 2025 zero-day review, drawing chiefly on commercial surveillance vendor activity against mobile targets, found that multiple exploit chains discovered in 2025 included three or more vulnerabilities, inflating the number of individual vulnerabilities required to achieve a single objective, as attackers adjusted to vendor mitigations that had closed off simpler routes. What frontier AI changes is who can build such a chain and how long it takes. The graveyard is no longer a collection of tolerable residual risk. Instead, it is a target inventory for adversaries with access to systems that can reason across vulnerability combinations autonomously, at machine speed, without the specialist skill that previously made such chains impractical.

In the post-Mythos world, underwriting cyber risk will require a shift in focus from validating the existence of security controls to understanding how organizations manage uncertainty and build resilience to manage dynamic threats and changing conditions.

Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

How Frontier AI Models Affect the Attacker Model

Cybercrime has experienced several waves of attacker industrialization.

Initial access brokers have created specialist marketplaces that enable threat actors to purchase access rather than obtain it themselves; malware-as-a-service providers have enabled threat actors without the technical capability to build their own tools, such as endpoint detection and response (EDR) killers, to buy them instead; and ransomware-as-a-service operators have transformed extortion into a scalable business model. Each development has changed attacker economics. Frontier AI acts on the same economics from a different direction: industrialization divided the labor among existing specialists, while frontier AI lowers the skill barrier itself.

Organizations have always managed vulnerabilities, misconfigurations, excessive trust relationships, and operational weaknesses. The significance of frontier AI models is not that they create new categories of cyber risk or even new attack techniques. Instead, they represent a structural change in how attackers can behave, changing the speed and scale at which weaknesses can be discovered, analyzed, and potentially exploited.

Historically, identifying a critical software vulnerability required considerable expertise, time, and persistence. Today, frontier AI models are demonstrating an ability to analyze large codebases, identify patterns, generate proof-of-concept exploits, and reason through complex attack paths with increasing effectiveness. Attackers are often unconstrained by the operational realities that organizations face. They do not manage production environments,

conduct testing cycles, seek change approvals, or balance security against business continuity.

Exploiting zero-day vulnerabilities was typically reserved for targeted operations, after which the same vulnerabilities were exploited by opportunistic criminals. Frontier AI should therefore be viewed less as a revolutionary change in attacker capability and more as another step in the ongoing reduction of attack costs and operational friction.

The implications of frontier AI extend beyond cybersecurity operations and reach directly into how cyber risk is assessed, priced, and managed.

The Mythos Moment: Why Cyber Insurance Should Pay Attention

Cyber insurers have never been able to predict when a significant vulnerability would emerge.

The discovery of a major zero-day has always represented a source of uncertainty: no insurer knows which vulnerability will be discovered next, who will discover it first, or whether it will be exploited before a fix becomes available. This uncertainty is not a flaw in cyber insurance. It is one of the reasons cyber insurance exists.

If vulnerability discovery becomes increasingly machine-assisted, the volume of newly identified weaknesses may rise significantly. The issue is not whether vulnerabilities exist. They always have. The issue is whether the rate at which they are discovered begins to accelerate faster than the ecosystem's ability to absorb them. How long that lasts is an open question and it matters for pricing. The current surge may be a one-time reckoning, machine-assisted analysis draining a finite backlog

Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

of accumulated defects in existing code. If so, disclosure volume spikes, then settles at a new baseline once the backlog clears. The alternative is that discovery capability compounds against code that continues to be written, in which case the elevated rate persists. The two paths imply different reserving assumptions, and the industry does not yet have the data to tell them apart. Insurers should treat the current volume as a transition signal rather than a permanent frequency. For insurers, this raises important questions about frequency assumptions, exposure modeling, and the future shape of cyber loss events.

Cyber insurance has always operated in an environment of uncertainty. What matters now is how that uncertainty is changing. The number of zero-day vulnerabilities may well decrease, met by a sharp increase in the number of Known Exploited Vulnerabilities (KEVs). For underwriters, that trade is not obviously bad. It converts unknown risk, which cannot be priced, into knowable risk, which can. A portfolio exposed to a thousand catalogued and remediable vulnerabilities is more tractable than one exposed to an unknown number of undiscovered ones. The burden shifts from uncertainty to capacity: the constraint becomes whether the insured can act on what is now visible.

The Traditional Underwriting Assumptions Being Challenged

Cyber insurance underwriting was built around human-scale offensive operations and AI-augmented enhancements.

Thematic conclusions emerging from projects such as Glasswing, ExploitGym, and MDASH are beginning to demonstrate a material shift in how threat actors can behave and accomplish

their objectives. The authors have identified three assumptions underlying traditional cyber insurance underwriting that are being challenged by recent technological innovations.

1. Discovery of vulnerabilities is stable and material zero-day impact is predictable

Mythos Preview found thousands of high-severity zero-days in every major operating system and browser within weeks. These were not marginal findings in obscure code but critical flaws in software the world runs on. The ExploitGym benchmark from Berkeley, Anthropic, OpenAI, and Google documented that frontier AI agents also routinely find vulnerabilities they were not looking for: in 69 of 226 successful exploits, agents achieved code execution through a vulnerability other than the one they were directed at, sometimes by auditing source code and performing dynamic fuzzing autonomously.

The discovery of zero-day vulnerabilities produced uncertainty that enabled growth in the cyber insurance market. No one knows when a zero-day will be discovered and exploited, who will be targeted, or how long it will remain active before a workable mitigation, patch or compensatory control, is identified. Cyber insurance helped organizations temper this uncertain risk.

However, that same uncertainty is the part insurers cannot price directly. Insurers cannot underwrite the existence of vulnerabilities nobody has found yet. Until a vulnerability is identified, assigned a CVE, and associated with detectable signatures, its probability and potential frequency cannot be quantified with confidence. However, insurers have access to exposure data that provides insights into

Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

the industries and sectors that rely on a common set of suppliers, platforms, and technologies. This allows insurers to model the consequences that a potential zero-day could have across a portfolio. This understanding enables insurers to recognize industries with heightened systemic exposure and reflect that concentrated risk through pricing or underwriting restrictions.

The change in the rate of vulnerability discovery poses both a threat and an advantage to insureds, depending on who has access to frontier AI models and their motivations for discovering vulnerabilities. With malicious intent, the likelihood of systemic supply chain catastrophes increases when actors discover new zero-day vulnerabilities at speed. This will be particularly true for open-source software dependencies and shared libraries. With altruistic intent, zero-day vulnerabilities can be recategorized as known vulnerabilities more quickly, helping organizations develop mitigation strategies before a material impact is realized.

2. A workable remediation relationship exists

The cyber ecosystem has historically relied on a workable relationship between discovery and remediation: a vulnerability is discovered, a vendor develops a fix, organizations deploy the patch, and exposure is reduced. This process has never been perfect. High-profile vulnerabilities such as Log4Shell and WannaCry, along with the exploitation of edge devices demonstrate that remediation has often lagged behind attacker activity. Even so, the system broadly functioned because discovery, disclosure, and remediation moved at roughly comparable speeds with the development and use of vulnerability exploits.

Frontier AI introduces pressure on that relationship. Discovery, exploit development, and weaponization may all accelerate, yet patch testing, operational approvals, deployment schedules, and business continuity requirements remain constrained by human and organizational realities. Large enterprises cannot safely patch everything immediately, regardless of how quickly vulnerabilities are identified.

If vulnerability disclosure pipelines begin receiving significantly greater volumes of findings, organizations may face

increasing challenges in prioritization and response. For insurers, the concern is not simply the existence of vulnerabilities but the widening gap between the rate at which exposure is identified and the rate at which it can realistically be reduced.

Even before AI-scale discovery, the volume of vulnerabilities affecting systems and software was outpacing the capacity to remediate them. End-of-life and out-of-support systems remain in production while new vulnerabilities continue to surface against them. Organizations have managed this backlog by prioritizing severity, using CVSS scores to rank criticality and the Cybersecurity and Infrastructure Security Agency's (CISA's) KEV catalog to identify what is being actively weaponized. That approach created the graveyard described above.

The cyber ecosystem has historically relied on a workable relationship between discovery and remediation.

Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

Frontier AI has changed what that graveyard means. Vulnerabilities that scored as medium or low individually and were treated as low priority accordingly can now be chained autonomously into exploit sequences that achieve full system compromise.

The ExploitGym benchmark, published in May 2026, documents AI agents routinely combining low- and medium-severity findings into complete attack chains in under two hours. The CVSS-based prioritization model was built on the assumption that severity scores approximate actual exploitability. That assumption no longer holds when an adversary can reason across a vulnerability inventory and identify chains that no single score predicts.

The second problem is temporal, and it compounds the first. The ExploitGym and MDASH results described above point the same direction: the window between AI-assisted discovery and a weaponized exploit is now measured in hours, and what arrives through it needs no credentials and no specialist skill to fire. Against that, the window between patch publication and enterprise deployment is still measured in weeks to months. The 2026 Verizon Data Breach Investigations Report (DBIR) puts the median time to fully resolve a known exploited vulnerability at 43 days, up from 32 the year before, with only 26% of KEV-listed vulnerabilities fully remediated in 2025, down from 38%. The 2025 edition had already put the median time from disclosure to mass exploitation of edge device vulnerabilities at zero days. Even with standard mitigations active, ASLR, stack canaries, and the V8 heap sandbox, ExploitGym shows frontier AI succeeding on a meaningful number of tasks.

The remediation window has not narrowed. It has inverted.

3. Insurer visibility is adequate to measure and transfer risk

Underwriters have never possessed perfect visibility into cyber risk. They rarely see internal architectures, security debt, configuration drift, development practices, or undiscovered vulnerabilities. Instead, they rely on observable indicators such as governance practices, questionnaire-based validation of implemented controls, external scanning, claims experience, and questions posed to the CISO or head of cybersecurity. These indicators have proven valuable because they provide a reasonable proxy for risk.

Today, a growing proportion of cyber risk may sit beneath the observable attack surface, particularly where AI systems identify exploit chains involving multiple low- and medium-severity weaknesses that traditionally would have been evaluated independently.

Cyber Risk Quantification (CRQ) models can take the above inputs to produce assumed loss magnitude and breach likelihood values, which must be interpreted correctly to model and manage risk through mitigation and transfer. All CRQ models can be thought of as both correct and incorrect. They are correct for the assumptions they consider and incorrect because cyber risk outcomes cannot be bounded by a single figure.

Mythos Preview discovered a remote code execution vulnerability in FreeBSD that provided unauthenticated root access that had existed undetected for 17 years. Microsoft's MDASH surfaced 16 previously undisclosed

Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

vulnerabilities in the Windows networking and authentication stack in a single Patch Tuesday cohort, four of them Critical remote code execution flaws. The FreeBSD flaw and the Patch Tuesday cohort were not surfaced by existing methods, underwriting reviews, or security testing activities. Vulnerability scanners would not have detected them because they were unknown vulnerabilities and therefore did not yet have known detectable characteristics.

Taken separately, each of these assumptions has been strained before. What is different now is that all three are under pressure at once, and they fail in the same direction. Discovery accelerates. Remediation does not. The insurer's view of the insured refreshes once a year. The result is a policy priced against a risk picture that can change materially between binding and renewal, with nothing in the process designed to notice that it has.

That is the new category. Not a new peril, and not a failure of underwriting practice, but exposure that moves inside the policy period while the instruments used to price it stand still..

Underwriting the Single Risk

The relationship between the insurer and the insureds has always contained information asymmetries.

Insured organizations will be able to take advantage of frontier AI models to support their own defensive operations or inherit the changes made by suppliers who update their own products. The insured's understanding of risk will improve, while the insurer's understanding may improve more slowly. Whether this becomes a material issue remains to be seen, but it

represents an important shift for a market built on the effective assessment of uncertainty.

If cyber risk becomes increasingly dynamic, insurers may need to reconsider which signals matter most. Understanding the existence of controls, external attack surface visibility, governance, and security culture all remain important. But future underwriting may increasingly prioritize indicators of adaptability and resilience, rather than patch velocity.

The following questions offer an additional underwriting signal focused on understanding the use of AI to augment defensive operations:

- Has the organization used AI-assisted discovery against its own environment (e.g., code repositories)? If so, does its AI security capability rely on a single-model scan or a multi-model agentic system with validation and proof stages?
- Can the AI model identify exploit paths rather than individual vulnerabilities?
- Does the organization use AI to augment vulnerability prioritization efforts? ?

The following questions are already considered in underwriting signals and should not be lost through the competitive landscape brought by soft-market conditions:

- How quickly can the organization contain a threat and recover when compromise occurs? ?
- Has the organization exercised its ability to respond to a cyber incident through a crisis simulation, one that recognizes there are often no right answers in a response, only less bad answers, and cyber incident investigations may often provide incomplete knowledge to support decision making?

Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

Repositioning the KEV Variable

KEVs serve an important purpose by identifying vulnerabilities known to be exploited in the wild, informing a prioritization mechanism for remediation.

In an environment where security teams are balancing remediation timelines, business governance, and finite resources against thousands of vulnerabilities, KEVs provide information on what should be prioritized first. However, KEVs do not provide certainty.

A KEV is a signal that exploitation has occurred somewhere. It does not confirm that every exposed organization will be or could be impacted.

As cyber insurance underwriting has matured, some insurers have incorporated KEV-related policy conditions that restrict coverage or apply sublimits if identified KEVs are not remediated within a defined timeframe. This approach risks treating KEV presence as a direct breach indicator rather than as a prioritization signal, and that risk may become more pronounced in a post-frontier AI world.

If frontier AI increases vulnerability discovery rates, the number of vulnerabilities entering disclosure pipelines is likely to increase. Over time, this may result in a larger volume of vulnerabilities progressing into KEV catalogs as exploitation activity follows disclosure. The practical consequence is that organizations may face a growing stream of KEV remediation requirements without any corresponding increase in their ability to test, approve, and deploy changes safely. In this situation, the existence of a KEV becomes less informative than an organization's ability to respond to it.

An organization that consistently identifies, prioritizes, and remediates KEVs is demonstrating

operational capability. It is showing that security, engineering, and operational teams can work together to reduce meaningful exposure under real-world constraints. Conversely, an organization

that repeatedly fails to address high-priority vulnerabilities may be demonstrating deeper challenges relating to governance, engineering practices, asset visibility, or security culture. Viewed through this lens, KEVs provide a proxy to measure organizational maturity. This distinction matters because underwriting should ultimately seek to understand how an organization responds when risk becomes visible. The presence of a vulnerability is often influenced by factors outside an organization's control, whereas the speed and effectiveness of its response are not.

The most valuable underwriting signal may therefore not be the existence of a KEV itself, but the demonstrated ability to remediate meaningful exposures consistently over time. Treating KEV detection as a policy limitation trigger risks conflating awareness of risk with the cause of loss. Treating KEV remediation performance as an indicator of operational maturity provides a more durable and potentially more predictive underwriting signal. The CISA BOD 26-04 Directive mentioned in Beyond Patch Velocity, Paper 1 in this series, also provides remediation timeline guidance for vulnerabilities aligned to CISA's new four-variables approach, ranging from 3 to 60 days, or remediation during the next system upgrade. These timelines can be used to guide the new underwriting signal.

The presence of a vulnerability is often influenced by factors outside an organization's control, whereas the speed and effectiveness of its response is not.

Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

Recognizing the Portfolio Risk

A single cyber event can generate losses across several insureds simultaneously. Frontier AI amplifies this reality by changing the speed, scale, frequency, and unpredictability with which accumulation events emerge.

Frontier AI amplifies this reality by changing the speed, scale, frequency, and unpredictability with which accumulation events emerge. For example, accelerated software supply chain compromises and faster exploitation of vulnerabilities could increase the likelihood of correlated losses across insurance portfolios within shorter time frames. For cyber insurers, the challenge here is less about understanding the probability of individual events and more about understanding how AI may alter systemic exposure across an entire book of business.

“No Regrets” Actions

While the impact of frontier AI on cyber insurance portfolios remains uncertain, several actions represent prudent “no regrets” investments that strengthen portfolio resilience regardless of how the threat landscape changes:

Insurers should revisit cyber catastrophe scenarios to test assumptions about the speed and scale of attack propagation. Scenarios that once assumed human-paced campaigns may underestimate the impact of AI-enabled attacks capable of exploiting common technologies or supply chain dependencies at unprecedented speed.

Insurers should place greater emphasis on exposure concentration analysis. Understanding concentrations of cloud providers, identity platforms, managed service providers, security technologies, and other shared digital dependencies is becoming increasingly important as these common technologies represent potential sources of correlated loss.

Insurers should continue to evolve aggregation modeling beyond individual technologies to capture interconnected digital ecosystems and third-party dependencies. Investments in cyber catastrophe modeling, threat intelligence, and network dependency analysis will help insurers better understand where systemic exposures exist within their portfolios.

Preparing for a More Uncertain Future

The longer-term response from insurers will depend on how frontier AI capabilities develop in practice and what is observed in the frequency and severity of claims:

- If AI substantially reduces the cost and expertise required to conduct sophisticated cyber attacks, insurers may need to tighten underwriting standards, revise pricing assumptions, and continue to actively manage concentrations across industries.
- If newly disclosed vulnerabilities can be weaponized almost immediately, traditional assumptions about the time available for organizations to patch critical systems may no longer hold. This would place greater emphasis on continuous exposure monitoring and faster portfolio risk assessment.

Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

Frontier AI may not substantially increase expected annual losses, but it could materially increase loss volatility and model uncertainty. Greater model uncertainty could influence capital allocation, catastrophe modeling, reinsurance purchasing, and portfolio diversification strategies, placing increased importance on stress testing and scenario analysis alongside traditional actuarial approaches. For insurers, this shifts the challenge from estimating average losses to managing uncertainty about extreme, correlated events.

Conclusion

Insurers occupy a unique position within the cyber risk ecosystem, observing claims, incidents, loss trends, and root causes across thousands of organizations, technologies, and sectors.

As the threat landscape becomes more dynamic, this portfolio-level perspective may become an increasingly valuable source of intelligence, helping policyholders understand emerging risks, benchmark resilience, and identify patterns that may not be visible from within a single organization.

Equally, cyber insurance has an important role to play in strengthening resilience. Many insurers already provide access to proactive security services, incident preparedness support, and specialist response capabilities that organizations would otherwise struggle to procure at scale. Insurers are increasingly positioned not only to transfer risk but also to help organizations reduce its impact when incidents occur.

The future of cyber insurance may therefore be defined by more than underwriting alone. It may be defined by the industry's ability to combine risk transfer, intelligence, and resilience into a single value proposition for policyholders.

The emergence of frontier AI does not invalidate the foundations of cyber insurance.

Cyber risk has always been uncertain, and insurers have always operated in a world where

vulnerabilities emerge unexpectedly, exploitation timelines are unpredictable, and perfect visibility is impossible. What may be changing is the scale of those challenges. As discovery accelerates, remediation struggles to keep pace, and exposure increasingly sits beneath the observable attack surface, insurers may need to rethink how they assess cyber risk.

The future of cyber underwriting is unlikely to be found in vulnerability counts alone. It will be found in understanding how organizations discover, manage, and respond to uncertainty, because the defining challenge of a post-frontier AI world is not that cyber risk may suddenly become unknowable. It is that the gap between what exists, what can be discovered, and what can be observed may be widening faster than ever before. Closing that gap also requires a conversation inside the insured organization that currently has no owner. Buying Resilience, Paper 3 in this series, takes up who is positioned to force it.

Cyber insurance has an important role to play in strengthening resilience.

Beyond Vulnerability Counts

Reimagining Cyber Insurance Underwriting in a Post-Frontier AI World

About the Authors

James Tuplin, is the Head of International Cyber underwriting at Mosaic Syndicate where he is responsible for the cyber underwriting teams based in London, Cologne, and Dubai. James has 23 years of underwriting experience specializing in cyber and technology E&O insurance for the last 20 years.

Jay Vinda, is the CISO and Cyber Risk Engineering Lead at Mosaic Insurance. He bridges the gap between technical security concepts, business objectives and risk management strategies. Through his work, he aims to contribute towards establishing cyber insurance as a powerful risk management partner, proactively helping organizations to protect systems and become resilient to emerging threats.

Sanith Wijesinghe, Dr. Sanith Wijesinghe is the Chief Engineer for MITRE's public sector. Over the past 10 years he has overseen MITRE's R&D initiatives in blockchain, including development of the AADAPT cyber threat framework for digital assets. Most recently he led a NIST cybersecurity workshop panel discussion on "Systemic Threats, Vulnerabilities, Interoperability and Other Shared Challenges" for digital ledger based technologies.

Chris Folk, is a Senior Cybersecurity Leader at MITRE, with 25 years in cybersecurity, national critical infrastructure protection, including eight years leading the CVE Program through its 2016 federation redesign. His work focuses on the governance layer of cyber risk: how disclosure, coordination, and interpretation reduce exposure risks.

